TIME SERIES FORECASTING OF AIRLIFT SUSTAINMENT CARGO DEMAND

GRADUATE RESEARCH PAPER

Daniel S. DeYoung, Major, USAF

AFIT/IMO/ENS/12-05

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

AFIT/IMO/ENS/12-05

TIME SERIES FORECASTING OF AIRLIFT SUSTAINMENT CARGO DEMAND

GRADUATE RESEARCH PAPER

Presented to the Faculty

Department of Operational Sciences

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Master of Science in Logistics

Daniel S. DeYoung, ME

Major, USAF

June 2012

AFIT/IMO/ENS/12-05

TIME SERIES FORECASTING OF AIRLIFT SUSTAINMENT CARGO DEMAND

Daniel S. DeYoung, ME
Major, USAF

Approved:

________//SIGNED//__________ ___25 May 12___
Jeffery D. Weir, Ph.D. (Chairman) Date

AFIT/IMO/ENS/12-05

Abstract

Forecasting demand for airlift of sustainment cargo is an important function for logistics planners. For the civil reserve air fleet participants (CRAF), having a useful long-range forecast enables them to make business decisions to maximize profit and manage their fleets. Because the DoD relies on CRAF for much of its steady-state and wartime surge requirements, it is important for these civilian enablers to stay financially healthy in what has become a difficult market. In addition to the CRAF airlines, DoD schedulers benefit from somewhat shorter-term forecasts of demand, as accurate forecasts help them allocate aircraft type and determine route frequency for airlift of sustainment cargo.

Time series forecasting is a method applied in many circumstances, to include forecasting of aviation service demand. It does not require the modeler to attribute causation, but rather uses historical data of a univariate series to predict future values. This paper applies a variety of time-series techniques to historical sustainment demand data from Iraq and Afghanistan AORs, ultimately choosing a single technique to develop into a prediction model for future demand in each AOR. The resulting models show excellent goodness-of-fit values and are successfully validated against a reserved portion of data.

Dedicated to my wife and our wonderful children

Acknowledgments

I sincerely appreciate the early guidance of Mr. Don Anderson, who helped connect me with this subject matter and was my conduit to the data. I especially would like to thank Jeff Weir, Ph.D., for his guidance and review of this effort along the way. Most importantly, thanks to my family for supporting my efforts to research and write.

Maj Daniel S. DeYoung

Table of Contents

Page

Page

List of Figures

Page

List of Tables

Page

List of Equations

Page

TIME SERIES FORECASTING OF AIRLIFT SUSTAINMENT CARGO DEMAND

I. Introduction

> "As the global war on terror[ism] continues, our forces are in distant countries fighting organized terrorists who seek to destroy our nation and destabilize the world. Military operations in these austere places are challenged by the need to deploy and supply troops over great distances. Airlift is a precious lifeline that keeps them fed and equipped, brings the wounded home, and eventually, brings our forces home."

Congressman Jim Saxton, 4 April 2005

General Issue and Problem Statement

The United States was at war in Iraq from March, 2003 until December, 2011, and has been at war in Afghanistan since October, 2001. These massive undertakings required the deployment and sustainment of tens of thousands of troops in two different areas of responsibility (AORs) and have highlighted the need for better tools to forecast sustainment demand. Figure 1 illustrates troop levels over time for these two conflicts. As with any military endeavor, but especially regarding those taking place far from home basing, logistics has been crucial to enabling operations. A great definition of military logistics crystallizes this role as follows:

> "logistics is the process of planning and executing the movement and sustainment of operating forces in the execution of military strategy and operations. It is the foundation of combat power – the bridge that connects the nation's industrial base to its operating forces". (Kress: 4)

A relatively large portion of the sustainment of troops in Southwest Asia is accomplished via air transportation, likely due to the remoteness of the region, particularly in the case of Afghanistan. For example, United States Transportation

Command (USTRANSCOM) through its organic air component Air Mobility Command (AMC) and its partnership with commercial carriers moved an estimated 20% of sustainment cargo to the Afghanistan AOR in 2010 (TRANSCOM/J3 Ops Update). AMC's wartime logistics mission is to bring supplies, material, and/or personnel to supply units already engaged in combat operations. (JP 3-17: I-2).

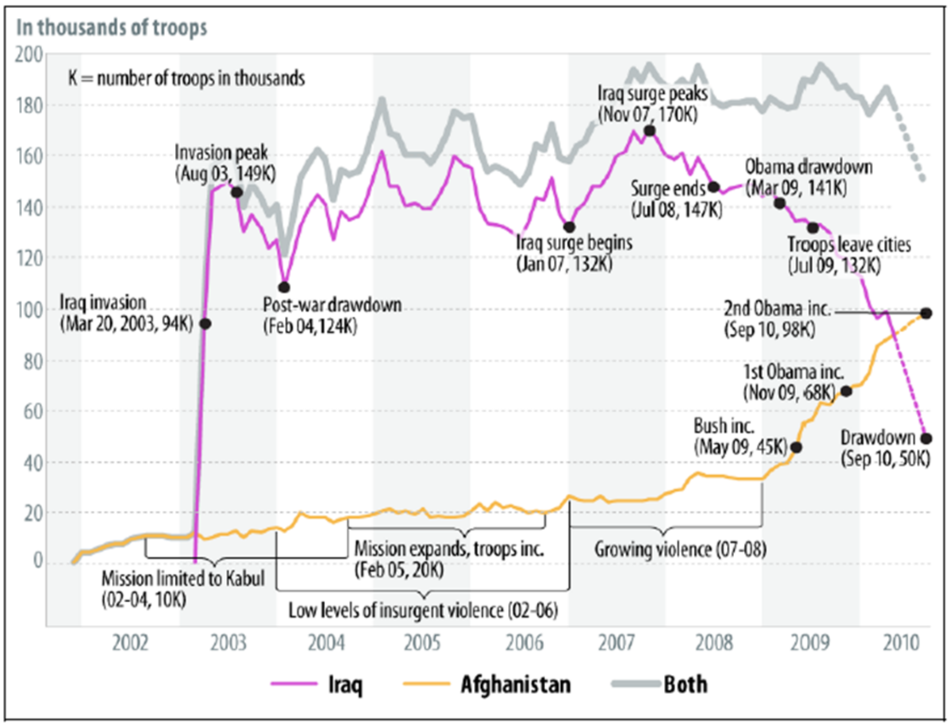

Source: DOD, *Boots on the Ground Reports to Congress*; CRS Report RL30588, *Afghanistan: Post-Taliban Governance, Security, and U.S. Policy*, by Kenneth Katzman; CRS Report RL31339, *Iraq: Post-Saddam Governance and Security*, by Kenneth Katzman; CRS Report RL34387, *Operation Iraqi Freedom: Strategies, Approaches, Results, and*

Figure 1: Boots on Ground

Unfortunately, forecasting levels of demand for airlift of sustainment cargo has been difficult. As Figure 1 illustrates, troop levels are constantly changing. When the boots-on-ground numbers change, sometimes rapidly as in the case of the Iraq and Afghanistan "surges", sustainment demand tends to change as well, although not always proportionately. This presents a dual challenge. First, DoD planners need to know the demand in order to optimize a mix of commercial and organic assets tailored to the demanded levels of cargo. In general, planners assign the Civil Reserve Air Fleet (CRAF) to channel missions, which are more predictable and less dangerous than other mission types. Consequently the CRAF "flies 72% of DOD channel cargo missions, freeing up the AMC organic fleet for deliveries to high threat areas" (Grismer: 7). The second challenge is to the CRAF participants, who need a long range outlook to optimize business decisions.

A brief discussion of commercial partners is warranted here as context to the importance of demand forecasting. The Civil Reserve Air Fleet was started just after World War II as an organized way to expand the nation's capacity for airlift in times of emergency. The DoD offers peacetime business to participating US-controlled airlines as an incentive for their commitment of no less than 30% of their qualifying passenger fleet and 15% of their qualifying cargo fleet (Graham: 29). In return, CRAF participants dedicate their designated fleet to the DoD for use when activated by the TRANSCOM commander (with approval of the Secretary of Defense) to one of three stages: minor regional crises (Stage 1), major theater war (Stage 2), or times of national mobilization (Stage 3) (Graham: 3). Unfortunately today's difficult economic environment and globalized airline industry have pushed our airlines to the edge of unprofitability. From 2001 to 2004 the industry lost $32 billion (Morrison: 1), returning to losses again with the recession beginning in 2008 (Census Bureau). For CRAF to continue to entice our commercial partners to provide valuable wartime service the program must be a profitable *and predictable* endeavor; effective long range scheduling is key to efficient, sustained profitable operations.

Research Objectives/Questions/Hypotheses

The primary research objective is to create a model useful in predicting sustainment cargo demand up to a year in advance to enable strategic planning for CRAF and organic assets. This long range forecasting capability is particularly useful for the CRAF partners, who "are best able to employ resources to support their primary commercial obligations and steady state DoD business when they can see all requirements in advance" (Grismer: 5). But it is also beneficial to the DoD which garners better participation from CRAF partners by forecasting the requirements (Grismer: 5). Referencing Figure 2, which is tailored to commercial passenger airlines but applicable nonetheless, the long-range forecast is an enabler for the CRAF in strategic planning, budgeting, and revenue management.

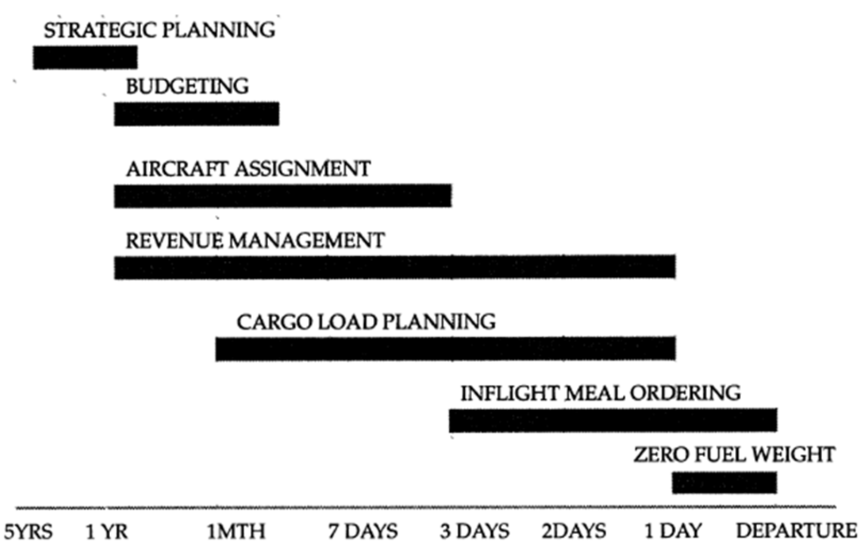

Figure 2: Forecasting Applications and Time-Frames Relative to Flight Departure

Source: Wickham

A secondary objective is for this model to be applied to the short-range (up to 3 month) scheduling of CRAF and organic assets. This shorter range capability would be useful to DoD planners who schedule channel missions and assign aircraft types to these missions. In Figure 2, these are the functions of aircraft assignment and cargo load planning.

Methodology

This research employs a time-series analysis of cargo data pulled from the Global Air Transportation Execution System (GATES) database. The analyzed data represent sustainment cargo whose aerial port of debarkation (APOD) resides in either of the two AORs or select other US Central Command (USCENTCOM) airfields. According to JP 3-17, "the vast majority of airlift sustainment will move on channel missions" (JP 3-17: I-6). The two types of channel mission, both captured in the data, are termed "requirements" and "frequency". Wartime requirements channels fly on schedules determined by traffic demand with set time parameters (time definite deliveries, or TDD), sacrificing efficiency for effectiveness if required. Frequency channels are more regular in order to give geographic combatant commanders a predictable way to move cargo and personnel. (JP 3-17: I-6) Therefore, to model sustainment cargo the researcher used only those data coded as "channel" missions from fiscal years 2005-2011. After limiting the data in this way, the total gross weight for the subject cargo, measured in pounds throughout this paper, was grouped into regular time-periods (seven-day and monthly) as required for time-series analysis. Gross weight was used because channel cargo scheduling and billing is generally done on a per-weight basis (JP 4-09: p. II-5). After applying and comparing a number of different time-series analysis techniques, the researcher used goodness-of-fit testing and validation techniques to form a model.

Assumptions/Limitations

The researcher assumes that any errors or omissions in the GATES data source are randomly distributed and minor in degree, meaning such errors would affect data throughout the time series in equally small ways. This assumption allows the resultant model function to be presented as representative of the historical cargo weight values and valid in predicting future values.

The researcher also assumes that approximations made in reducing the data source to pertinent entries do not significantly affect the efficacy of the model. For example, as previously mentioned the researcher equated "channel" cargo with "sustainment" cargo, which is true the "vast majority" (but not all) of the time. Appendix A, Table 1 lists the airfields whose cargo demand is represented in the model.

While the researcher was not adversely limited by the quantity of available data, the nature of this data did shape the methodology. The data contained useful information to track a pallet's weight, volume, route and dates of travel, and type of aircraft used, but it lacked information on significant independent variables (such as unit type requesting cargo) that might allow for a meaningful regression analysis of demand. Given the available data, time-series forecasting was the most appropriate tool.

II. Literature Review

Chapter Overview

This research required a working knowledge of CRAF as a backdrop for the motivation behind long-range forecasting. There is a vast amount written on many aspects of the CRAF program; this research reviewed articles which helped explain the need, from the commercial partners' point of view, for a more reliable and longer-range forecast of DoD business. The selected literature coupled a description of the airline industry business environment with a look at the risks to doing business with the DoD.

The heart of the literature review focuses on demand forecasting. There is much written about demand for things from widgets to electricity to airline passengers and cargo, covering both commercial and government concerns, and ranging from strategic (long-range, big picture) forecasts to tactical (short-range, detailed) forecasts. The pertinent literature addresses analytical techniques used in forecasting demand.

Finally, the researcher performed a search for literature on airlift sustainment cargo demand forecasting. While much has been written on cargo demand forecasting in general, very little was discovered dealing specifically with the subject of this paper, forecasting demand for wartime airlift of sustainment cargo. A brief review of existing literature is nevertheless presented.

CRAF and the Airline Industry

Since its inception CRAF has only been activated twice. During Operation Desert Storm (ODS) long range cargo and passenger aircraft were activated to Stage I from Aug 1990 to Jan 1991, then to Stage II in Feb 1991. During Operation Iraqi Freedom (OIF) the CRAF saw Stage I activation for passengers in Jun 2003 (Arthur: 3). During activation CRAF assets fly missions in support of the DoD and are compensated at negotiated ton-mile and passenger-mile rates, but control of scheduling CRAF assets is given over to Air Mobility Command. Support for the DoD comes at the expense of commercial revenues and can mean a loss of market share where competitors are not as heavily invested in CRAF. As an example, the DoD did not meet CRAF program participation goals for the two years subsequent to Operation Desert Shield/Storm. This was due to significant concerns by CRAF participant CEOs regarding long-term business loss to competitors while their aircraft were committed to military operations (Arthur: 3).

As a result the DoD had to improve the incentives for CRAF participation, one of a few times such a move has been required. Beginning in 1995, participating carriers could only gain access to "peacetime" government business by committing assets to the CRAF. Table 1 illustrates overall industry contribution to the CRAF program, while Table 2 shows specific airline requirements to access government customers.

Table 1: Industry Commitment to CRAF

Aircraft Committed to the Civil Reserve Air Fleet in June 2007, by Segment

	Cargo Aircraft			Passenger Aircraft		
	Stage I	Stage II	Stage III	Stage I	Stage II	Stage III
International Segment						
Long-range section	31	79	300	42	147	690
Short-range section	n.a.	12	12	n.a.	13	270
National Segment						
Domestic services section	n.a.	0	0	n.a.	24	37
Alaskan section	n.a.	4	4	n.a.	0	0
Aeromedical Evacuation Segment	n.a.	0	0	n.a.	25	50

Source: Congressional Budget Office using information from the United States Transportation Command.

Notes: Internal divisions of the Civil Reserve Air Fleet (segments and sections) reflect the unique capabilities required of committed aircraft.

n.a. = not applicable.

Source: Congressional Budget Office "Issues Regarding the Current and Future Use of the Civil Reserve Air Fleet," 2007.

Table 2: CRAF-Linked Government Business

Business Category	Annual Business ($ millions)	CRAF-Linkage Provisions	Comment
Business that Requires a Minimum Fleet Commitment			
GSA City Pairs Program (Individually ticketed DoD, other government, and government contractor personnel.)	$ 560 million	Commitment: 30 % of long-range fleet	Total City Pairs eligible revenue is about $1 billion; Due to waivers and incomplete enforcement, about 40% of eligible revenue goes to non-CRAF carriers
Domestic Charter (Full-planeload domestic passenger charters.)	$ 62 million	Commitment: 30 % of long-range fleet	
Express Cargo – Domestic (Domestic, small parcel, office to office shipment for DoD other government agencies, and cost-reimbursable government contractors)	$ 98 million	Commitment: 25 % of long-range fleet; 30% for both domestic and international	
World-Wide Express Cargo (International, small parcel, office to office shipment for DoD other government agencies, and cost-reimbursable government contractors)	$ 35 million	Commitment: 25 % of long-range fleet; 30% for both domestic and international	
Business that is Allocated in Proportion to Commitments (Mobilization Value Points)			
Category A Cargo (Palletized cargo, less than full planeload, pick up and drop off at military depot.)	$ 55 million	Business allocated in proportion to MV points; Minimum Commitment: 15 % of long-range fleet	
AMC Passenger Charter (International full-aircraft passenger charters.)	$ 300 million	Business allocated in proportion to MV points; Minimum Commitment: 30 % of long-range fleet	"Fixed-buy" charters are specified for each contract year; "Expansion buys" meet short-notice requirements.
AMC Cargo Charter (International full-aircraft cargo charters.)	$ 180 million	Business allocated in proportion to MV points; Minimum Commitment: 15 % of long-range fleet	"Fixed-buy" charters are specified for each contract year; "Expansion buys" meet short-notice requirements. Cargo Charter Business will be displaced by organic C-17s

Source: Institute for Defense Analysis, "Sustaining the Civil Reserve Air Fleet (CRAF) Program,": 30)

Access to roughly 60% of this business requires an airline to commit at least 30% of its long-haul capacity to CRAF. Access to the remaining 40% of government business is achieved through the mobilization value (MV) point system, which allows teams of carriers to form and compete for government business. MV points are earned according to the number and type of aircraft committed to CRAF, and government market share is earned in proportion to MV points (Graham: 29-30).

The peacetime (non-activated) business is divided into "fixed buys" and "expansion buys". Fixed buys are purchased at the outset of a fiscal year and constitute a contract between the DoD and the airlines for a number of guaranteed payments for particular routes flown. These payments are made at the time service is provided unless the DoD under-tasks relative to the fixed buy, in which case the difference is dispensed at the end of the fiscal year. "Expansion buys" are anything tasked above and beyond the fixed buy due to unforeseen or underestimated requirements. Figure 3 shows the relative proportion of fixed buy to expansion buy through 2007, with estimates thereafter. Although the expansion buys have become much greater in terms of total revenue, it is the fixed buy that offers a more effective incentive. According to the CBO, "those guaranteed payments are a particularly attractive incentive to carriers to participate in the Civil Reserve Air Fleet because they can count on those funds in formulating their annual business plans" (Arthur: 4).

As Figure 3 indicates, the proportion of fixed buy to expansion buy is already too low, meaning the fixed buy incentive leaves much room for optimization. In addition, the DoD is predicting a drop in its need for commercial airlift support after the war in

Afghanistan ceases, which in turn will lead to a drop in the potential fixed buy incentive.

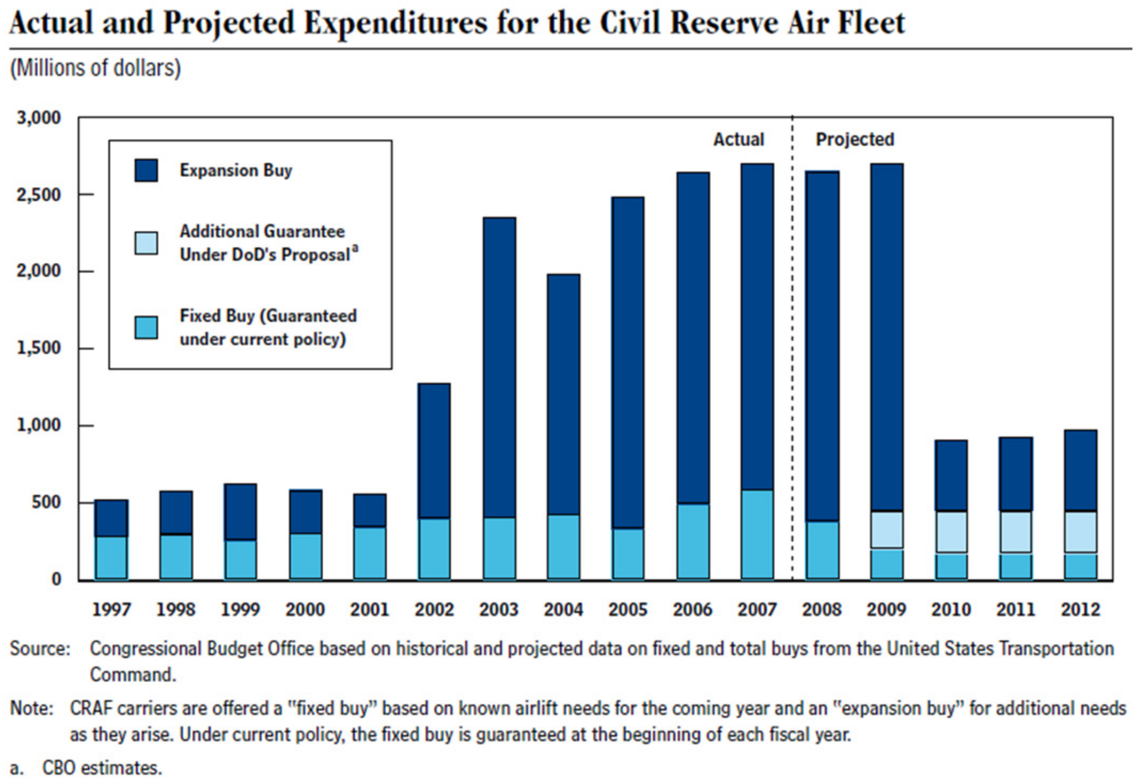

Figure 3: Fixed vs. Expansion Buys

Taken together, these threats to the fixed-buy incentive spurred another proposal to adjust the CRAF incentive: an increase to the proportion of business guaranteed in the fixed buy. This FY08 DoD proposal would take a rolling average of the total (fixed + expansion buy) business over the previous five years, omitting any unusually high outliers. The DoD would then be able to offer up to 80% of this rolling average in the fixed buy (Arthur: 2). The success of this method is dependent on the accuracy of the rolling average method, which is also a type of time series forecast. Although it was never implemented, it may resurface in a future attempt at higher fidelity forecasts. In

Appendix B, the researcher proposes that the model put forth in this paper is an improvement to the moving average proposal.

Demand Forecasting

Demand forecasting is an integral part of many operational and financial concerns, and as such a wide variety of research can be found in the literature addressing the topic. For example, Stoimenova, et. al. investigated using time-series models to forecast demand for electricity, concluding their forecasts would be useful in predicting periods of peak demand, capacity and maintenance planning, and investment decisions. Chen, et. al. investigated demand planning approaches to forecasting the need for safety stock in business, arguing that "with the globalization of demand–supply networks and the desire for a more integrated operation plan, demand planning is now one of the greatest challenges facing manufacturers" (Chen, et. al.: 1). Velonius used linear regression to predict supply and demand levels for freight tankers, thus creating an overall model to forecast tanker freight rates. The preceding examples show application in the utility, manufacturing, and transportation sectors of the economy.

While good information can be gleaned from an analysis of forecasting as applied in these various circumstances, the researcher narrowed the literature to forecasting phenomena in the realm of aviation so as to more closely follow the subject of this research. An excellent work discussing many forecasting techniques applicable to air passenger demand is "New Directions for Forecasting Air Travel Passenger Demand" by Garvett and Taneja. Although published in 1974, the paper discusses many forecasting techniques still relied on today. As the authors put it, despite recent developments in new *methods* for forecasting (the Box-Jenkins, or ARIMA, method used in this paper was new at the time) "we are still a long way from possessing *tools* which provide our decision-makers with more effective, that is, more useful, accurate and timely information" (Garvett: 1, italics added). As stated, this research seeks to provide exactly such a tool.

Garvett and Taneja distinguish between qualitative and quantitative approaches to forecasting. The qualitative techniques are used when little historical data exists, or otherwise if the trend of the historical data is expected to change. Therefore the qualitative techniques are better suited to big-picture, strategic type forecasts, dealing more with what is possible than in the timing of specific events. Examples of qualitative forecasting include S-curve, Delphi studies, and morphological analysis (Garvett: 3-5).

Quantitative techniques are more widely used. They depend on a body of historical data from which to extend an historical trend. Quantitative techniques can be divided into time-series methods and causal methods. Time-series methods are generally analyses of one-parameter data taken over regularly-spaced time intervals, whose trends can be extended to predict future events. Examples include moving averages, spectral analysis, and the Box-Jenkins method. Causal methods seek to quantify the effect of individual (independent) variables on the output in question. These methods are less time-dependent and more event-dependent. Examples include regression models, Bayesian analysis, Markov chains, input-output analysis, simulation methods and control theory models (Garvett: 3-5).

More will be discussed on the specifics of time-series techniques in the section III, Methodology. However, it is fair to say that the Garvett and Taneja paper addresses various techniques and gives strengths and weaknesses thereof as they apply to aviation demand in general. They do not propose a single model to solve a given problem in particular. They conclude that air travel demand prediction models post-1974 would likely be structural in nature, meaning the underlying causal factors determining demand would be modeled. Identifying these casual factors, such as through a regression model, enables an accurate forecast of the effects of future changes (Garvett: 159). Yet the authors also make the following two points regarding time-series analysis which are apropos to this research effort:

- First, while other methods may appear more appealing based on theoretical grounds, data may not be available to justify their use.
- Second, a model's simplicity is in the mind of its user. Box-Jenkins methods and spectral analysis can hardly be classified as "simple and rough" and their use has considerably increased the validity of trend-extrapolation (Garvett: 19)

A final point they make regards adaptive forecasting. Static forecasts simply fit a smoothing equation to data from a set timeframe, determining constants and smoothing parameters which best fit that data. Future known values are not fed back into the model to enable adaption of the model to changing conditions. Garvett and Taneja suggest use of adaptive forecasting, in which the model is updated with new data. Two types of adaptive forecasting exist; in the first, the model itself changes through addition or subtraction of variables. Of this type, they argue “it is rare that the addition of new variables will correct misspecification errors in a weak model” (Garvett: 20). Instead the authors advocate the second type, in which the basic structure of the model itself remains intact, but the parameters and constants of the model are updated to best fit the new data.

The Box-Jenkins approach is an example of a model designed to be updated as new data is input, and thus is an adaptive forecasting technique (Garvett:29).

Another informative piece is Wickham's 1995 doctoral thesis, "Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation". In it he writes that forecasting short-term demand is crucial to the airline industry, enabling them to make tactical decisions and plan their transportation "supply" to meet market demand.

Table 3: Summary of Wickham's Models

Model #	Classification	Method
1	Time Series	Simple mean of final bookings
2	Time Series	Exponential smoothing of final bkd (α=0.2)
2b	Time Series	Exponential smoothing of final bkd (α=0.4)
3	Regression	Fbkd = f(bkd_t): same booking class
3b	Regression	Fbkd = f(bkd_t): different booking class
4	Classical Pickup	Simple mean of total pickup
5	Classical Pickup	Exponential smoothing of total pickup (α=0.2)
5b	Classical Pickup	Exponential smoothing of incremental pickup (α=0.4)
6	Advanced Pickup	Simple mean of incremental pickup
7	Advanced Pickup	Exponential smoothing of incremental pickup (α=0.2)
7b	Advanced Pickup	Exponential smoothing of incremental pickup (α=0.4)

Source: Wickham

Wickham compared variations on seven different models to determine their effectiveness in predicting passenger demand on a particular 18-week set of data (Table 3). The goal was to predict future demand while also testing the effects of limiting the quantity of historical data and extending the forecast period. The time series approaches used simple mean and exponential smoothing (without seasonality). The regression models assumed a relationship between passenger demand on the day in question and passenger demand seven days prior. The many "pickup" models related the quantity demanded on the day in question to the quantity already booked on day *n*, plus the expected number of "pickups" between day *n* and the day of takeoff; in essence, these models are hybrids incorporating both previous techniques. The classical pickup models only use data from flights already departed to average the number of expected "pickups" *n* days out. The advanced pickup models also include bookings made for flights that have yet to take off, thereby increasing the data pool by adding more recent data (Wickham: 45-54). His conclusions showed that the pickup models did better than both time series and regression models, with the advanced pickup models achieving the best results.

Applying Wickham's study to this research, the first observation is that the time series models used were of the simplest design and did not incorporate seasonality. Wickham's own research highlights that "seasonal variation occurs quite naturally in the demand for air travel", with holidays causing spikes and the whole first quarter of the calendar year constituting the "off-season" (Wickham: 22). However, the data he selected for analysis did not span a complete calendar year, and so he was unable to incorporate seasonal indices into the data set. Wickham assumes this omission will have little effect, but that assumption leaves much room for error and is a weakness in the study (Wickham: 67). In addition, the pickup models don't immediately translate to the forecasting of cargo demand. The concept of a hybrid approach, however, might have merit here and will be recommended as an area for future study.

A final relevant article presented under demand forecasting is “Air Cargo Demand Prediction”, by Totamane et. al. As the title indicates, this 2009 article deals directly with predicting air cargo demand where the previous two did not. However, while the preceding articles used time-series analysis in their analysis, Totamane uses a technique called a weighted majority algorithm. The purpose of their research was to use demand prediction to help tailor a commercial airline’s flight schedule and cargo capacity plan, in order to optimize both load factors and delivery success rates (Totamane).

While this technique is not informative of the method used in this paper, it is nevertheless briefly included in the literature review to reveal the type of forecasting that has been applied to air cargo in particular. The weighted majority algorithm, as applied by Totamane, allows for an airline to select certain predictors of cargo demand such as day-of-the-week cargo data, projected regional economic growth, quarterly cargo supply averages, special occasion (holiday) historical data, and naïve predictors. The airline then relies on its booking agents’ skills at interpreting these predictors, giving each agent a weight according to their record of accuracy. The weighted majority algorithm then assembles a demand forecast by integrating the weighted estimates of each predictor by the many agents for any given day (Totamane).

Sustainment Cargo Airlift Demand Forecasting

The researcher found only one published piece dealing directly with the subject of this paper. Billion and Regan’s 1966 presentation “A Methodology for Accurately Predicting Demand for Airlift of Military Cargo to Overseas Destinations” is a relevant, if dated, research effort. The authors were commissioned by the DoD to measure the economic benefit gained by using the newly acquired C-5A strategic airlifter to help satisfy peacetime cargo movements (Billion). They were critical of time-series extrapolations as applied to air cargo demand in the commercial sector, arguing such methods were not successful in terms of forecasting actual cargo weights delivered. Figure 4, taken from their paper, seems to support this claim. Thus, instead of applying time-series forecasting to the prediction of military resupply the authors created a new model based on the quantity and frequency of demand for specific commodities to overseas destinations. In effect, the model was a tool to help decide which resupply commodities might be more economically flown by air than transported by sea to resupply DoD bases. The relevant costs considered were loss and damage, packaging, pipeline, in transit warehousing, inland line haul, and line haul costs.

Billions and Regan's research made extensive use of "approximately 3.8 million commodities and millions of shipments recorded on magnetic automatic data processing

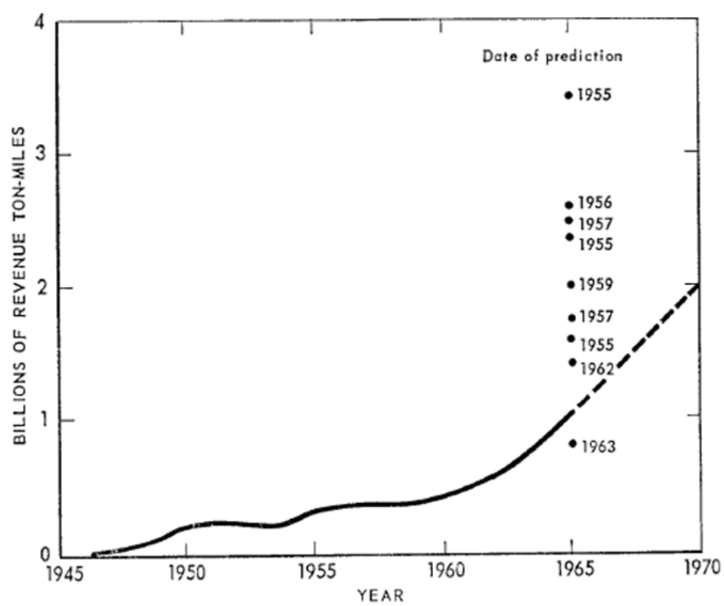

Figure 4: Air Freight Forecasts Compared with Tonnage Generated (US Domestic)

tapes” (Billions). Their model seems like a good predictor of demand for static conditions. That is, their model depended on calculating elasticity of demand for specific commodities to the DoD’s overseas footprint as it existed in 1966 or as it was projected to be in the coming years. As they explained, “the data base for the work is necessarily empirical, but demand for certain classes of commodities can be correlated with present force structures. By basing calculations on Joint Chiefs of Staff estimates of force structures in the 1970’s, the total demand can be projected…” (Billions). The data is a good fit for the known force structure, but is dependent on JCS predictions of future overseas force structure for its forecasting accuracy. This pegs the accuracy of their model to that of JCS force structure projections, leaving them open to the type of error seen in Figure 4 above.

In addition, as Garvett advised, the better model is the one whose structure is maintained but whose parameters are updated as new data are available. The Billions and Regan model would have to be fully reconstructed as old commodities become obsolete and new ones are added, as well as whenever overseas force structures undergo meaningful change. Referring back to Figure 1, Boots on Ground, this would be labor intensive indeed if applied to the wars in Southwest Asia.

Finally, their critique of time-series forecasting predates the newer, more powerful time-series techniques used in this work. Box and Jenkins, for example, published their work on ARIMA techniques in 1970. The many smoothing techniques described later in Methodology weren’t proposed until 1957 by Holt, and it wasn’t until 1963 that the forms commonly used today were put forward by Brown. The predictions graphed in Figure 4 were mostly made in the 1950s during a period punctuated with growth spurts in the domestic airline industry. The models underlying those predictions were not presented, but since the forecasting community has since moved on to more effective methods that also better model trending and seasonality it is a safe assumption that the authors' critique needs to be revisited.

Summary

This literature review presented background information on the CRAF program to explain why the program participants desire a long range business forecast from the DoD. The review then covered demand forecasting, revealing its widespread use and giving examples of various forecasting techniques as they are applied to specific uses. The researcher then focused on demand forecasting for air industry concerns, including forecasts for both passenger and cargo demands, and discussed the one example which closely mirrors this research effort. The research reveals the following points which set the stage for this research effort:

1. Previous relevant time-series applications did not use the more powerful techniques available today
2. Previous time-series applications may have been misapplied with respect to accurately modeling trends and inclusion of seasonality
3. An adaptive model, whose structure remains intact as its parameters are updated by new data, is preferable to a static model, or to one whose structure must change to adapt

Therefore, this research applies the more powerful time-series methods to create an adaptive tool for decision makers in industry and the DoD alike.

III. Methodology

Chapter Overview

The DoD maintains a number of bases within USCENTCOM, where both Iraq and Afghanistan AORs reside. While this list of bases sometimes changes, most of the significant locations remain. Appendix A lists the locations included in this research effort. Shipment histories to and from these fields (and all other DoD fields) are recorded in the GATES database. Understandably, this database is very large and in the context of data analysis can be unwieldy. This chapter explains how the researcher scoped the data to the pertinent entries and describes the data itself, detailing which fields are significant to the research and what information can be gleaned from them. Thereafter is presented a discussion of the different time-series methods employed in analyzing the data. Finally, this chapter discusses model formulation to include a description of the JMP software used for data analysis.

Scope and Data Description

An important first step for this research was scoping the data. The data sources are GATES databases for FY2005-FY2011, provided as Microsoft Access files by Air Mobility Command. The researcher exported these database files into Microsoft Excel for data reduction. The typical fiscal year of GATES cargo data contains approximately 300,000 entries, where an entry is the movement of a particular piece of cargo between two locations. Fortunately, the GATES database is a wealth of information regarding the shipment of each cargo piece. There are 30 fields (database columns) which in part: identify the cargo (mission and cargo ID numbers), describe the cargo (volume, height and weight, rolling stock or pallet), describe the locations of transportation nodes (APOE, or aerial port of embarkation, and APOD, or aerial port of debarkation), list departure and arrival dates and times, and give the final destination of each line of cargo (pallet APOD). This last field is significant, as the landing airfield for an aircraft flying a particular entry in the GATES database is often not the final destination of that cargo, i.e. the mission APOD is not necessarily the pallet APOD. This happens anytime cargo makes an enroute stop, either to be transloaded or for that aircraft to accumulate more cargo.

As alluded to, this research approximates sustainment cargo as equivalent to channel cargo. Singling out channel cargo is made possible by decoding the mission ID. Air Mobility Command's "MAF Mission ID Encode/Decode Procedures" holds the key to this coding (Table 4).

Table 4: Decoding Channel Missions

Second Character		Third Character	
B	Channel Cargo	A	Requirements Channel, Atlantic Region
J	Positioning to first onload	B	Frequency Channel, Atlantic Region
K	Channel PAX	C	Frequency Channel, Pacific Region
L	Aeromedical Evacuation (AE)	P	Requirements Channel, Pacific Region
Q	Channel Mixed (PAX and cargo)	E/Q	CPX Channel missions as assigned by AMC/A3TX
V	Depositioning from offload to new mission or home station	J/R/U W/Y/Z	Channel missions supporting Contingency Operations as assigned by 618 TACC/XOP

The researcher applied the following process using Excel macros to focus the data:

1. The first step was to limit the database to channel cargo, meaning those entries whose mission ID had a "B" for the second character.

2. Next the Pacific missions were removed (C and P values from the "third character" column in Table 4).

3. Then all entries whose pallet APODs were NOT on the scoped list in Appendix A were removed, including entries for which pallet APOD was blank (generally some small number were).
4. Finally, entries with duplicate pallet IDs had to be reduced to one representative entry. This step addressed the problem outlined above, where pallets might be counted multiple times due to enroute stops, artificially increasing the measured demand signal.

These steps reduced the data by roughly 75%, leaving the researcher with data representing sustainment cargo to fields within USCENTCOM, predominantly in the Iraq and Afghanistan AORs. It should be noted that not all airfields retained in the model fall into those two AORs. The other retained USCENTCOM locations were, in the researcher's estimation, significant hubs for operations enabling activities in Iraq and Afghanistan, or otherwise nevertheless contributors to the demand signal for strategic airlift of sustainment cargo into USCENTCOM.

An interesting point became evident during data reduction: disaggregating the data led to better model results. This concept is supported in the research as well: "Aggregation during the model construction phase of the analysis will cloud the underlying behavioral relationships and will result in loss of information. It is always desirable to estimate a model at the disaggregate level" (Ben-Akiva: 62). For this research the data is disaggregated into subsets: Iraq data into one subset, and Afghanistan plus all other modeled CENTCOM fields into the other (hereafter referred to as Afghanistan+). The Iraq data is modeled separately for two reasons. Firstly, operations in Iraq have almost entirely come to a close, so modeling Iraq demand separately describes its contribution to the historical data while allowing it to be removed from the predictive tool going forward. Secondly, the Iraq and Afghanistan AORs can conceivably have distinct demand signals, each with their own seasonality, thus separate models are appropriate.

The quantity of available data is sufficient to create a forecasting model. The significant fields in the GATES data are the dates and locations of each entry, which allowed the researcher to filter the data by location, sum the gross weights (demand) and group these totals by date. The remaining GATES columns do not add significant information pertaining to forecasting of demand. Given data primarily describing gross weight delivered over time to particular locations, time-series forecasting methods are a natural fit.

Forecasting Techniques

This section describes the time-series forecasting methods applied in this research, with the addition of naïve and moving average forecasts as background. They are presented in order of increasing functionality.

Naïve Forecasts and Simple Moving Averages

Naïve forecasts predict the next value of a function to be equal to the last observed value. While simplistic, there are situations when this approach can achieve adequate results. An added benefit is its independence from the need for large data sets and computing power. The equation of a naïve forecast is shown below:

$$\hat{Y}_{t+k} = Y_t \qquad (1)$$

Where:
$\hat{Y}$ = forecast value

The moving average model takes the average value of the series over the previous *k* observations. A *k* value of 1 would be a naïve forecast, whereas a *k* value which includes all series observations would simply be the series mean. Averaging mitigates the rapid changes which can occur with naïve forecasting (Duke.edu). An example equation follows:

$$\hat{Y}_{t+1} = \frac{Y_t + Y_{t-1} + ...Y_{t-k}}{k} \qquad (2)$$

Where:

$\hat{Y}$ = forecast value

Simple Exponential Smoothing

The simple moving average model above only incorporates the previous *k* observations, which it weights equally. The idea behind exponential smoothing is to incorporate all series values in forming the model, while giving the more recent values more weight as it is assumed they are more relevant to predicting the next series value. The technique is called "exponential" because the weighting is done using a geometric series. Substituting previous values for *Y* on the right side of Equation 3 would introduce these terms. The choice of α is very important to optimizing the model. High α values put more weight on the most recent observation, while low α values put more weight on historical values (psu.edu). An advantage of this model over the moving average is that it is adaptive: α can be continually updated with new data to optimize the model, usually by minimizing the mean squared error (duke.edu). The model takes the form:

$$\hat{Y}_{t+1} = Y_t - (1-\alpha)e_t \tag{3}$$

Where:
$\hat{Y}$ = forecast value
α = mean term smoothing constant between 0 and 1
e = the prediction error

Double (Brown's) Exponential Smoothing

Double exponential smoothing can model a trend in the data, indicating by a series mean which varies over time such that it is described by a sloped line. This model combines a smoothed value of the mean term with a smoothed value of the first differences of the mean term and includes the smoothing constant. The model is adaptive, addresses trending, but does not address seasonality. Its equations follow:

$$L_t = \alpha Y_t + (1-\alpha)L_{t-1} \tag{4.1}$$

$$T_t = \alpha(L_t - L_{t-1}) + (1-\alpha)T_{t-1} \tag{4.2}$$

$$\hat{Y}_t(k) = L_t + ((k-1) + \frac{1}{\alpha})T_t \tag{4.3}$$

Where:
L = smoothed value of mean term

T = smoothed value of estimated trend
α = mean term smoothing constant between 0 and 1
$\hat{Y}$ = forecast value

Linear (Holt's) Exponential Smoothing

Linear exponential smoothing is another approach to modeling trend in the data, but differs because it includes a second smoothing constant which can also be optimized. Its equations follow:

$$L_t = \alpha Y_t + (1-\alpha)(L_{t-1} + T_{t-1}) \tag{5.1}$$

$$T_t = \gamma(L_t - L_{t-1}) + (1-\gamma)T_{t-1} \tag{5.2}$$

$$\hat{Y}_t(k) = L_t + kT_t \tag{5.3}$$

Where:
L = smoothed value of mean term
T = smoothed value of estimated trend
α = smoothing constant between 0 and 1
γ = trend smoothing constant between 0 and 1
$\hat{Y}$ = forecast value

Seasonal Exponential Smoothing

By now it is obvious that each successive smoothing model builds on the original simple exponential smoothing model by adding some new aspect of the data to improve accuracy. The seasonal exponential smoothing model adds a seasonality component, although it does not include trending. That is, the data show a tendency toward some periodicity, but the mean tends to remain constant. The modeling equations follow:

$$L_t = \alpha(Y_t - S_{t-p}) + (1-\alpha)L_{t-1} \tag{6.1}$$

$$S_t = \delta(Y_t - L_t) + (1-\delta)S_{t-p} \tag{6.2}$$

$$\hat{Y}_t(k) = L_t + S_{t-p+k} \tag{6.3}$$

Where:
L = smoothed value of mean term
S = smoothed value of seasonality term
α = smoothing constant between 0 and 1
δ = seasonality smoothing constant between 0 and 1
$\hat{Y}$ = forecast value

Winters Method

The final method presented under the smoothing models is Winters method. It incorporates the smoothed mean term, smoothed trending term, and smoothed seasonality term, each with their own smoothing constant which can be optimized to fit new data. The equations which form this model are:

$$L_t = \alpha(Y_t - S_{t-p}) + (1-\alpha)(L_{t-1} + T_{t-1}) \tag{7.1}$$

$$T_t = \gamma(L_t - L_{t-1}) + (1-\gamma)T_{t-1} \tag{7.2}$$

$$S_t = \delta(Y_t - L_t) + (1-\delta)S_{t-p} \tag{7.3}$$

$$\hat{Y}_t(k) = L_t + kT_1 + S_{t-p+k} \tag{7.4}$$

Where:
L = smoothed value of mean term
T = smoothed value of the trending term
S = smoothed value of seasonality term
α = smoothing constant between 0 and 1
γ = trending smoothing constant between 0 and 1
δ = seasonality smoothing constant between 0 and 1

$\hat{Y}$ = forecast value

ARIMA (Box-Jenkins) Method

ARIMA stands for "Auto-Regressive Integrated Moving Average". This method, formulated by Box and Jenkins in 1970, is actually a more general form of every smoothing equation already presented. Stated another way, for each of the preceding smoothing model functions, there exists a corresponding and equivalent ARIMA function—see Table 5. The auto-regressive (AR) term is a linear regression of the value of a series observation against one or more previous series values. Thus the AR term is fitted to previous values of the series. The moving average (MA) term is a linear regression of the current series value on the noise evident in preceding observations in the series. Thus the MA term is fitted to previous values of noise in the series—actually, to all previous values of noise in the series. The term "integrated" (I) applies when the series must first be 'differenced' before applying the AR and/or MA analysis (NIST.gov).

Differencing is done to satisfy a prerequisite for ARIMA application, namely that the series is stationary with respect to its mean. If the series mean is not already stationary, this can be accomplished by subtracting the series mean from the series value at each point. If differencing one time (first differencing) does not achieve stationarity, the second difference can be taken. As a general rule, higher orders of differencing are not advisable (Duke.edu). Whatever degree of differencing that is applied during model creation must be removed, or *integrated*, before expressing the resultant time-series function.

Table 5: Smoothing Model and ARIMA Equivalence

Smoothing Model	ARIMA Equivalent
Simple Exponential Smoothing	ARIMA (0,1,1)
Double Exponential Smoothing	ARIMA (0,2,2)
Linear Exponential Smoothing	ARIMA (0,2,2)
Seasonal Exponential Smoothing	ARIMA $(0,1, p+1)(0,1,0)_p$
Winters Method	ARIMA $(0,1, p+1)(0,1,0)_p$

The ARIMA model is identified using the convention ARIMA(p,d,q), where p indicates the order of the AR term, d indicates the order of differencing (the I term), and q indicates the order of the MA term. The challenge in applying this method is in determining these values. Duke University's website is extremely useful in providing guidance toward forming the model—their rules for identifying ARIMA models are reprinted in Appendix D. Following these rules, the researcher tested the ARIMA(0,1,1) model.

Seasonal ARIMA

Seasonal ARIMA adds a seasonal term whose convention is similar to the original term. The new convention is (S)ARIMA(p,d,q)(P,D,Q), where some texts add the S and some do not. P indicates the order of the AR term within seasonality, D indicates the order of differencing for the seasonality term, and Q is the order of the MA seasonality term. P, D, and Q can be different values than p, d, and q. Thus, one possible example might be ARIMA(0,1,1)(1,0,0). However, the rules in Appendix D led the researcher to test an ARIMA(0,1,1)(0,1,1) model for this data set.

Model Development

The researcher used JMP 8.0 by SAS to analyze the data after completing the Excel filtering and grouping. JMP was used to plot the original series, perform time-series analysis using each of the above smoothing and ARIMA methods, and compare the results. For each method JMP displays a model summary listing various goodness-of-fit measurements, parameter estimates, a forecast graph depicting the model superimposed on the data, and information on the residuals. JMP also displays a chart comparing and ranking each method applied. The next section, Analysis and Results, presents these JMP outputs and decides on the best model to forecast demand for airlift of sustainment cargo.

Summary

The researcher's data is sourced from the GATES database, FY 2005-2011. Because each year of data is roughly 300,000 lines of cargo movement, most unrelated to sustainment demand in the USCENTCOM region, the researcher used Excel coding to perform four steps to filter and group the data. The scoped data, roughly 25% of the original data set, was then analyzed in JMP using the exponential smoothing and ARIMA methods described above. It is also noted that each smoothing equation can be described by an ARIMA equation. The next section will attest to this fact, as the best model fit and its ARIMA equivalent are almost equivalent in their model summary scores.

IV. Analysis and Results

Chapter Overview

This section presents and analyzes the JMP model-fitting results for a one-year forecast for the Iraq and Afghanistan+ data subsets, using the methods just discussed. Time series analysis requires data to be grouped into regular intervals; this research looks at both weekly and monthly groupings. The models, built using FY 05-10 data, are compared on the basis of the goodness-of-fit measures explained and presented below. The best-fit model for each grouping is then validated using FY11 data, and their forecasts and residuals are presented graphically. Validation techniques are also explained in this section. Appendix I builds a model for Kandahar AB, Afghanistan, to demonstrate the three-month forecast tool for use by DoD planners.

Goodness-of-Fit Measures

JMP outputs a number of goodness-of-fit measures. The RSquare, MAPE, and MAE values give an indication of how well the model fits the historical data. The -2LogLikelihood, Akaike's 'A' Information Criteria, and Schwartz's Bayesian Criterion values incorporate a penalty for model complexity (number of parameters), and are useful in ranking multiple models against one another. It is important to note that none of these measures indicate how well a model will predict future values; holding data in reserve for model validation serves this purpose. The goodness-of-fit measures are briefly defined below as background for the JMP outputs that follow in this section.

-2LogLikelihood

The likelihood function describes the probability density function for a model's parameter value(s) given realized outcomes of the model's random variable(s) (Rutgers.edu). A model can be described as the family of probability density functions that attempt to describe some phenomenon, with a model's parameter influencing the shape of that density function. In the inverse, when we already have known data values, we want to solve for the model parameter that shapes the density function such that the known outcome is the most probable outcome (Myung: 3). In the context of time series forecasting, with historical data being fit, the likelihood function gives the probability that the historical data is resultant from a particular model parameter value. This likelihood function can be maximized, that is, it can be solved for the parameter that best fits the data. Maximizing the function is often made mathematically easier by first taking the log of the function (Myung: 4). Therefore, -2LogLikelihood is minus two times the natural log of the likelihood function using the best-fit parameter. Smaller values indicate a better fit (JMP Support). A weakness of the likelihood function is that it does not penalize for higher numbers of parameters; judging a model by -2LogLikelihood alone might result in over-parameterization—improving the model's fit by inflating the goodness-of-fit measures with extraneous parameters.

Akaike's 'A' Information Criterion

Akaike's Information Criterion, or AIC, alters the -2LogLikelihood as follows:

$$\text{AIC} = -2\log\text{likelihood} + 2\text{m} \tag{8}$$

Where:

m = # parameters in model

The AIC method adds a penalty for each additional parameter, thus discouraging over-parameterization. In model selection, lower AIC values are better.

Schwarz's Bayesian Criterion

Schwarz's Bayesian Criterion, or SBC, is mathematically defined as:

$$\text{SBC} = -2\log\text{likelihood} + \text{m}\log\text{N} \tag{9}$$

Where:

m = # parameters in model
N = # of observations

Like AIC, SBC adds a penalty for each additional parameter. Lower SBC values indicate a better model.

RSquare

RSquare measures the percent of variation in a series accounted for by a given model. It is computed by:

$$\text{RSquare} = 1 - \frac{\text{SSE}}{\text{SST}} \tag{10}$$

Where:

$\text{SSE} = \sum_{i=1}^{N} (\text{y}_{\text{i}} - \hat{y}_i)^2$

$\hat{y}_i$ = the forecast series values

$\text{SST} = \sum_{i=1}^{N} (\text{y}_{\text{i}} - \overline{y}_i)^2$

$\overline{y}_i$ = the mean of series values

It can be seen that if sum of squares error (SSE) is larger than the sum of squares total (SST), RSquare will be negative, indicating a poor fit. On the other hand, an RSquare approaching a value of one would indicate a very small SSE in relation to SST, indicating a good model fit.

RSquare Adjusted

The adjusted RSquare value attempts to correct over-parameterization of a model by offsetting the RSquare score as more parameters are added. Without adjusting RSquare, we could expect the RSquare value to increase as the number of model parameters increase, although this "better" score likely does not translate to better predictive value. Therefore, adjusted RSquare is better at judging model accuracy while also deterring overuse of parameters. It is given by the following equation:

$$\text{RSquare Adjusted} = 1 - \left[\frac{(N-1)}{(N-m)} (1 - RSquare) \right] \tag{11}$$

Where:

m = # parameters in model
N = # of observations

MAE

The mean absolute error, or MAE, expresses the series mean of the absolute error:

$$MAE = \frac{1}{N}\sum_{i=1}^{N}\left|y_i - \hat{y}_i\right| \qquad (12)$$

Where:

$\hat{y}_i$ = the forecast series values
N = # of observations

Expressing the error in absolute terms ensures that errors of opposite sign do not have a cancelling effect. The unit of error is tied to the unit of the series data, making MAE difficult to use in comparing data sets of differing units.

MAPE

The mean absolute percent error, or MAPE, expresses the series mean of the absolute value of the fit error as a percent:

$$MAPE = \frac{100}{N}\sum_{i=1}^{N}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \qquad (13)$$

Where:

$\hat{y}_i$ = the forecast series values
N = # of observations

Expressing the error as a percent lends the advantage of comparing data sets expressed in different units (e.g. demand expressed in pounds vs. kg). MAPE is a simple and effective measurement of fit accuracy, yet it has the potential drawbacks of being undefined for any series value equal to zero, and being unbounded in terms of upper value. For the GATES data sample in this research, none of the aggregate demand values are zero.

AIC and SBC, as described above, arrive at parameter values by determining maximum likelihood estimation (MLE) for a parameter. RSquare attempts to minimize the sum of the squared errors, and therefore applies a least-squared error (LSE) method. These two approaches often yield different results, and when this occurs the MLE approach should be preferred (Myung: 6). Indeed, JMP solves for the model parameters using the log likelihood, an MLE approach upon which AIC and SBC are computed. The model goodness-of-fit outputs that follow are assessed according to their AIC and SBC values, MAPE values, and RSquare.

Validation Measures

Whereas goodness-of-fit measures assess how well a model matches the historical data, validation determines how well the model "performs" on future data. Model validation techniques include plotting the actual versus predicted "future" data, and computing values for the root mean square error of validation (RMSEv) and reduction of error (RE) statistics. These measures are defined below.

RMSEv

The root mean square error of validation is the mean size of the prediction error during the validation period. Comparing RMSEv to the RMSEc (calibration period RMSE, using historical data) gives an indication of model performance. RMSEv will generally be somewhat larger than RMSEc, as the latter is intentionally fit to the historical data. If the difference between the two is “small”, the model is considered validated. Determining what constitutes “small” is a somewhat subjective matter (Arizona.edu: 4).

$$RMSEv = \sqrt{\frac{\sum_{i=1}^{N_v} (y_i - \hat{y}_i)^2}{N_v}} \tag{14}$$

Where:

$\hat{y}_i$ = the forecast series values
N_v = # of observations in validation period

Reduction of Error

The reduction of error statistic is analogous to the RSquare goodness-of-fit measurement. It can range in values from negative infinity to one. There are two ways to interpret this statistic. First, with negative values the model is considered a poor predictor. Positive values indicate some degree of predictive quality, and a value of one indicates perfect prediction. The second interpretation is to compare the RE value to the

RSquare value; if they are "close" in value the model is considered validated. Again there is a degree of subjectivity inherent in the validation process (Arizona.edu: 4).

$$RE = 1 - \frac{SSE_v}{SSE_{ref}} = 1 - \frac{\sum_{i=1}^{N_v} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{N_v} (y_i - \bar{y}_c)^2} \tag{15}$$

Where:

$\hat{y}_i$ = the forecast series values
$\bar{y}_c$ = the mean series prediction over calibration period
N_v = # of observations in validation period

Weekly Grouping Models

Iraq Model

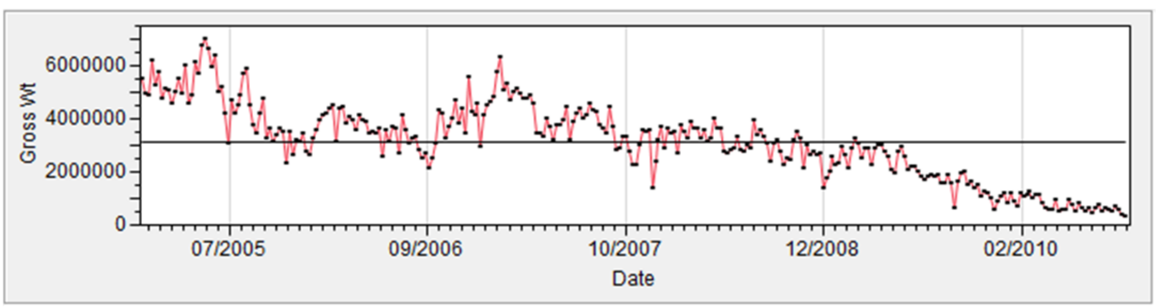

Figure 5: Time Series of Weekly Iraq Demand, Jan 2005-Sep 2010

Weekly Iraq demand from January 2005 through September 2010 is shown graphically in Figure 5, an output of JMP's time series analysis. It is evident that the demand from Iraq peaked early, leveled off through March of 2009, and has declined since. A seasonal pattern is also present, with a 52-week cycle.

The JMP model comparison output for the exponential smoothing and ARIMA techniques is presented in Figure 6. The AIC and SBC rankings for seasonal exponential smoothing and the seasonal ARIMA models are 1 and 2, although they disagree on which is the better of the two.

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Simple Exponential Smoothing	298	2.873e+11	8738.6044	8742.3048	0.861	8736.6044	5	5	15.612514	397905.96
Double (Brown) Exponential Smoothing	297	3.177e+11	8742.7032	8746.4003	0.843	8740.7032	7	6	16.013379	425848.35
Linear (Holt) Exponential Smoothing	296	2.881e+11	8718.3366	8725.7308	0.855	8714.3366	4	4	15.435702	404451.52
Seasonal Exponential Smoothing	245	2.853e+11	7265.2304	7272.2492	0.777	7261.2304	2	1	21.694237	476973.39
Winters Method (Additive)	244	3.283e+11	7278.5924	7289.1206	0.774	7272.5924	3	3	21.370458	474752.91
IMA(1, 1)	297	2.871e+11	8739.4121	8746.8130	0.862	8735.4121	6	7	15.239668	397037.55
Seasonal ARIMA(0, 1, 1)(0, 1, 1)52	244	3.047e+11	7264.1362	7274.6644	0.781	7258.1362	1	2	21.306261	470161.21

Figure 6: Model Comparison, Iraq Weekly Demand

The RSquare for these two methods are actually lower, and MAPE values higher than those of the remaining methods; the highest RSquare values are for the simple exponential smoothing and ARIMA models. Thus the goodness-of-fit criteria give conflicting indications for model selection, but a look at the forecast graphs in Appendix E helps to influence model choice. It can be seen that the non-seasonal methods give a linear forecast which doesn't very well approximate the seasonal influence on demand; yet due to the drawdown in Iraq, the seasonal influence on demand is largely diminished. In addition, aside from the simple exponential smoothing model, the rest predict negative values of demand.

Model Selection

On the basis of the model comparison chart in Figure 6 and the forecast graphs in Appendix E, the researcher chooses the simple exponential smoothing model for further analysis and validation. The final steps in forming the model are a check of the residual plot and a validation of the model on a portion of the data.

The residuals chart for the selected model is shown in Figure 7. There is a narrowing tendency in the residuals plot attributable to the twilighting of the conflict, but overall the graph shows a normally distributed plot of residuals with a mean of zero.

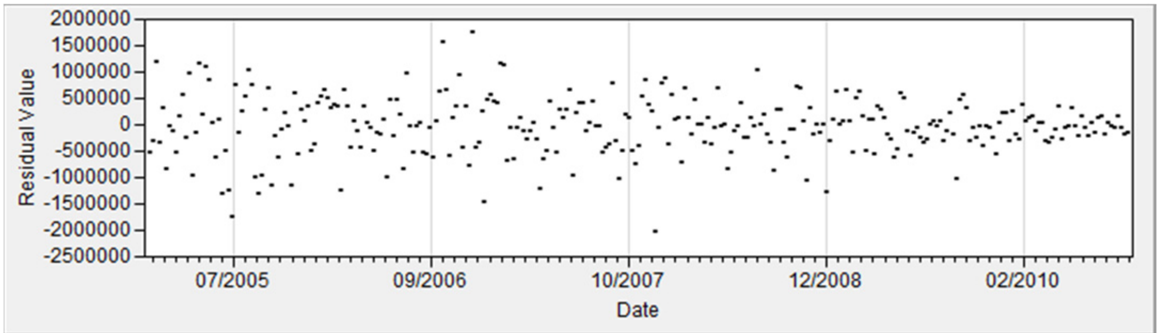

Figure 7: Residuals Plot for SES, Weekly Iraq Demand

Finally, The resulting fit and forecast graph is given in Figure 8.

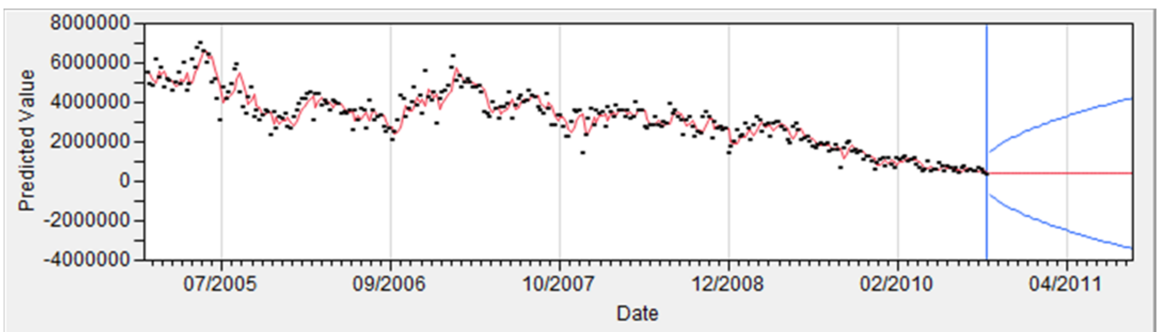

Figure 8: SES Forecast, Weekly Iraq Demand

Focusing on the forecasted region, and in keeping with the primary research objective, the following graph shows the predicted versus actual values for the series over a one year forecast period.

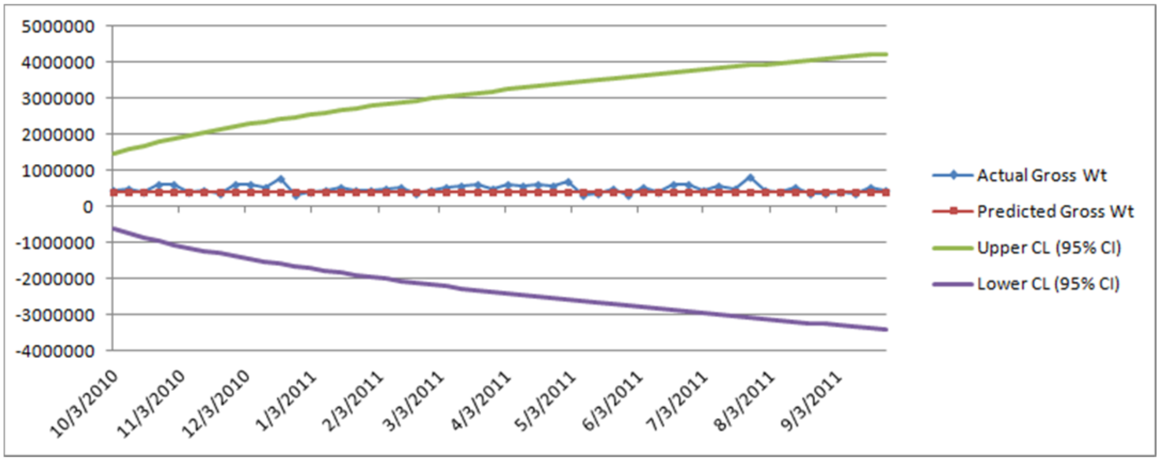

Figure 9: SES 1 -Year Forecast, Weekly Iraq Demand

Model Validation

Table 6: Validation Statistics for Simple Exponential Smoothing Model

RMSEv=	135628.6712	RE=	0.997378924
RMSEc=	535394.5339	R2=	0.861
% Difference	74.67	% Difference	15.84

Table 6 shows the validation statistics for the chosen simple exponential smoothing model. In this case the RMSE values are not roughly equivalent, which actually supports previous observations about the Iraq data. The demand became much easier to predict with the drawdown, which is why the residuals plot (Figure 7) narrowed as well. However, the validation period shows an explainable *increased* accuracy, not a failure in the forecast model. The RE and R2 values are closer in value, with RE almost equal to one. Therefore, the model is considered validated.

Model Equation

The model equation is given by Equation 3. The smoothing constant as computed by JMP is: α = 0.4897, with a t-ratio of 8.32. Ratios greater than two are significant.

Afghanistan+ Model

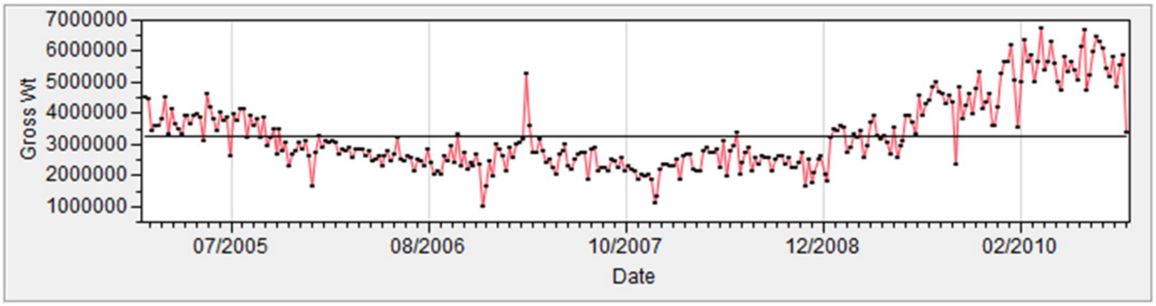

Figure 10: Time Series of Weekly Afghanistan+ Demand, Jan 2005 – Sep 2010

Weekly Afghanistan+ demand from January 2005 through September 2010 is shown graphically in Figure 10. The demand curve saw a reduction from the beginning of 2005 through 2009, likely due to sharing resources with Iraq operations. Post-Iraq demand is definitely higher, corresponding with the increase in boots-on-ground in Afghanistan from Figure 1. A seasonal pattern is also present, with a 52-week cycle.

The JMP model comparison output for the exponential smoothing and ARIMA techniques is presented in Figure 11. The AIC and SBC rankings agree on seasonal exponential smoothing as the best model. The RSquare for this method, while not the highest, is grouped closely with the best RSquare models. Yet the MAPE values are higher than most of the remaining methods.

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Simple Exponential Smoothing	298	2.869e+11	8738.5669	8742.2674	0.786	8736.5669	6	6	13.208608	397702.94
Double (Brown) Exponential Smoothing	297	2.979e+11	8725.8367	8729.5338	0.774	8723.8367	5	5	13.600100	410181.05
Linear (Holt) Exponential Smoothing	296	2.9e+11	8719.8018	8727.1960	0.780	8715.8018	4	4	13.350943	406132.97
Seasonal Exponential Smoothing	245	2.938e+11	7241.1540	7248.1727	0.800	7237.154	1	1	14.580904	435571.32
Winters Method (Additive)	244	2.945e+11	7241.4842	7252.0124	0.801	7235.4842	2	2	14.535087	431404.76
IMA(1, 1)	297	2.878e+11	8740.4573	8747.8581	0.786	8736.4573	7	7	13.244667	397424.64
Seasonal ARIMA(0, 1, 1)(0, 1, 1)52	244	3.027e+11	7242.9078	7253.4360	0.800	7236.9078	3	3	14.612166	432403.65

Figure 11: Model Comparison, Weekly Afghanistan+ Demand

Again the goodness-of-fit criteria give conflicting indications for model selection, so the forecast graphs in Appendix F help to influence model choice. It can be seen that the non-seasonal methods give a linear forecast which doesn't very well approximate the seasonal influence on demand. The seasonal ARIMA model's forecast increases more than should be expected for sustainment operations in what is apparently a new steady state (plateau) in demand.

Model Selection

On the basis of the model comparison chart in Figure 12 and the forecast graph in Appendix F, the researcher chooses the seasonal exponential smoothing model for further analysis and validation. Aside from its MAPE score, which is nevertheless competitive, its goodness-of-fit criteria are better than the competition and its forecast graph makes sense given current conditions. The final steps in forming the model are a check of the residual plot and a validation of the model on a portion of the data.

The residuals chart for the selected model is shown in Figure 12. The residual plot looks normally distributed, with constant variance and a mean of zero, as required.

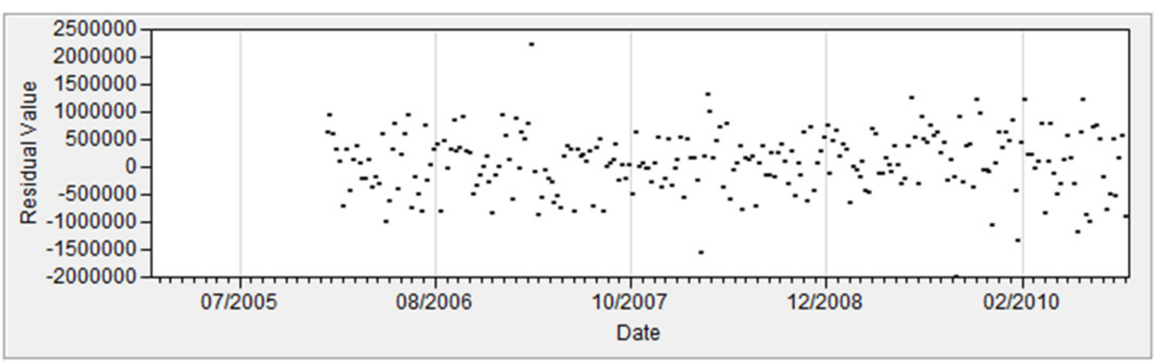

Figure 12: Residuals for Seasonal Exponential Smoothing Model, Weekly Afghanistan+ Demand

Finally, the model fit and forecast graph is given in Figure 13.

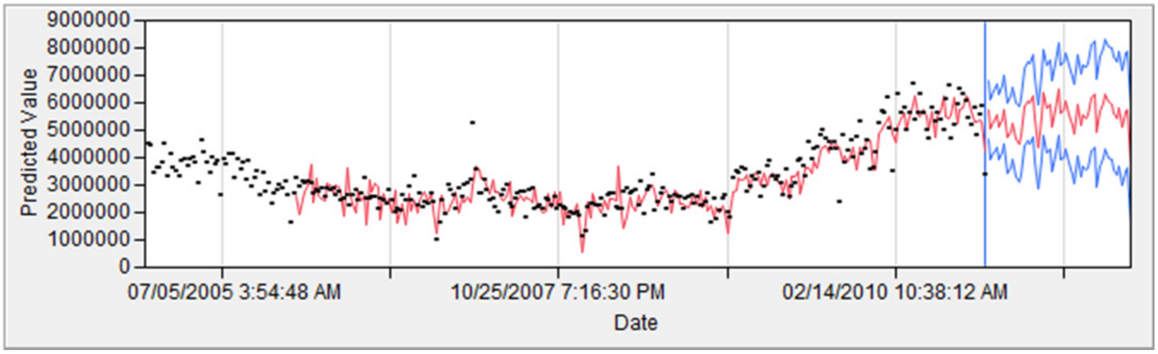

Figure 13: Seasonal Exponential Smoothing Forecast, Weekly Afghanistan+ Demand

Focusing on the forecasted region, the Figure 14 shows the predicted versus actual values for the series over one year. The confidence interval is graphed and the actual data fall within it.

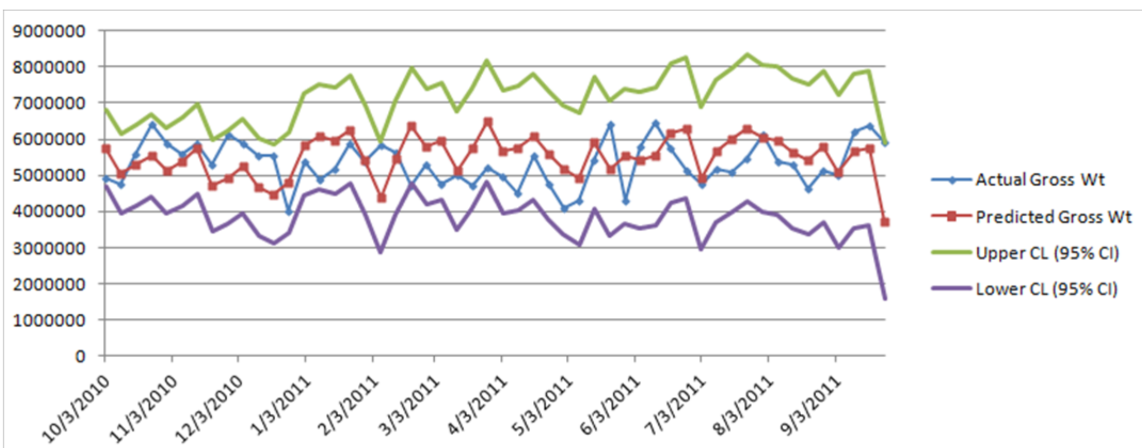

Figure 14: Seasonal Exponential Smoothing 1-Year Forecast, Weekly Afghanistan+ Demand

Model Validation

Table 7: Validation Statistics for Seasonal Exponential Smoothing Model

RMSEv=	849554.9571	RE=	0.857512351
RMSEc=	554928.7643	R2=	0.898
% Difference	53.09	% Difference	4.51

Table 7 shows the validation statistics for the chosen seasonal exponential smoothing model. It is arguable whether the RMSE values are roughly equivalent, but a clearer result is given by the strong parity between RE and R2 values. Given the highly positive RE value and its close approximation of R2, the model is considered validated.

Model Equation

The model equation is given by Equations 6.1 – 6.3. The smoothing constants as computed by JMP are:

α = .2486, with a t-ratio of 6.44.
δ = .7728, with a t-ratio of 7.50.

Again, ratios greater than two are considered statistically significant.

Monthly Grouping Models

Iraq Model

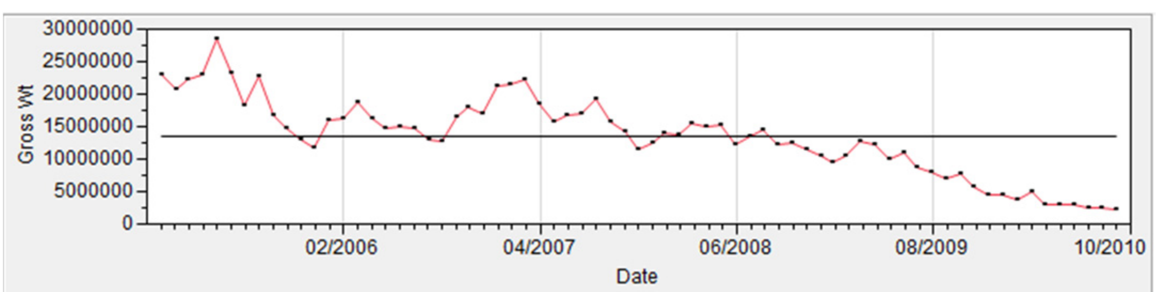

Figure 15: Time Series of Monthly Iraq Demand, Jan 2005 – Sep 2010

Monthly Iraq demand from January 2005 through September 2010 is shown graphically in Figure 15. As with the weekly aggregation, the demand from Iraq peaked early in the series, leveled off through the beginning of 2009, and has declined since. A slight seasonal pattern is also present, with a 12 month cycle.

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Simple Exponential Smoothing	67	4.924e+12	2181.3012	2183.5207	0.859	2179.3012	6	6	13.529292	1694362.0
Double (Brown) Exponential Smoothing	66	6.011e+12	2164.1909	2166.3956	0.818	2162.1909	5	5	13.990596	1851127.8
Linear (Holt) Exponential Smoothing	65	4.943e+12	2155.0301	2159.4395	0.839	2151.0301	4	4	13.167442	1749292.4
Seasonal Exponential Smoothing	54	4.76e+12	1810.7558	1814.8065	0.749	1806.7558	1	1	19.120519	2010147.0
Winters Method (Additive)	53	4.85e+12	1812.7558	1818.8318	0.749	1806.7558	2	2	19.120519	2010147.0
IMA(1, 1)	66	4.866e+12	2181.4816	2185.9206	0.863	2177.4816	7	7	12.728247	1675981.0
Seasonal ARIMA(0, 1, 1)(0, 1, 1)12	53	5.17e+12	1812.9005	1818.9766	0.750	1806.9005	3	3	19.532948	2008426.4

Figure 16: Model Comparison, Monthly Iraq Demand

The JMP model comparison output for the exponential smoothing and ARIMA techniques is presented in Figure 16. The AIC and SBC rank seasonal exponential smoothing as best. As with weekly demand, the RSquare for the seasonal methods are actually lower, and MAPE values higher than those of the remaining methods. The simple exponential smoothing model scores the best RSquare. Thus the goodness-of-fit criteria give conflicting indications for model selection.

The trending models' trend parameters do not assess as significantly different than the null. The seasonal models' seasonal component has a weak t-ratio for the seasonal exponential smoothing and Winters methods (.48), although the SARIMA model shows a t-ratio of 2.18, which is just high enough to be considered significant.

The forecast graphs in Appendix G are not significant differentiators. The problem of negative demand prediction is not as pronounced as with the weekly grouping, but still the simple exponential smoothing model is the only one that does not suffer from it.

Model Selection

On the basis of the model comparison chart in Figure 16 and the forecast graphs in Appendix G, the researcher chose the simple exponential smoothing and SARIMA models for further analysis and validation. The one-year forecasts for each method follow:

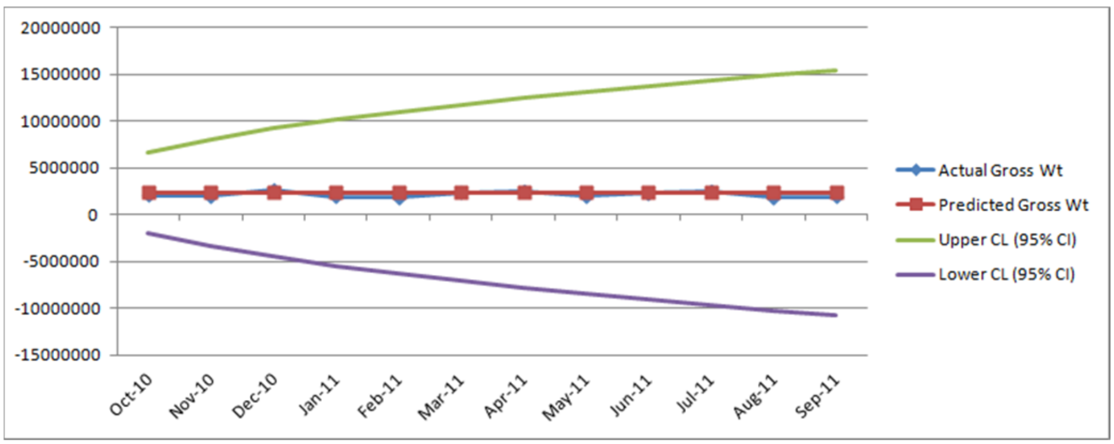

Figure 17: SES 1-Year Forecast, Monthly Iraq Demand

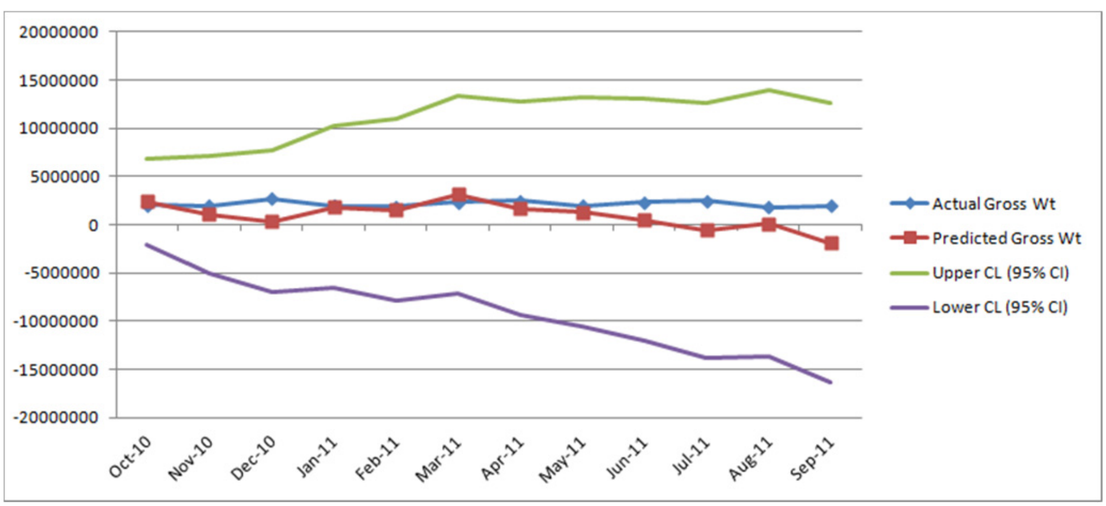

Figure 18: ARIMA(0,1,1)(0,1,1) 1-Year Forecast, Monthly Iraq Demand

The demand pattern in Iraq certainly changed over time due to changes in the underlying conditions in Iraq. With the drawdown complete, the best model for predicting future demand is the simple exponential smoothing model. The researcher's choice of model would likely have been different for example in 2007-2008, when the demand fluctuation and seasonality were stronger. In that case, the SARIMA method would likely have provided a better model.

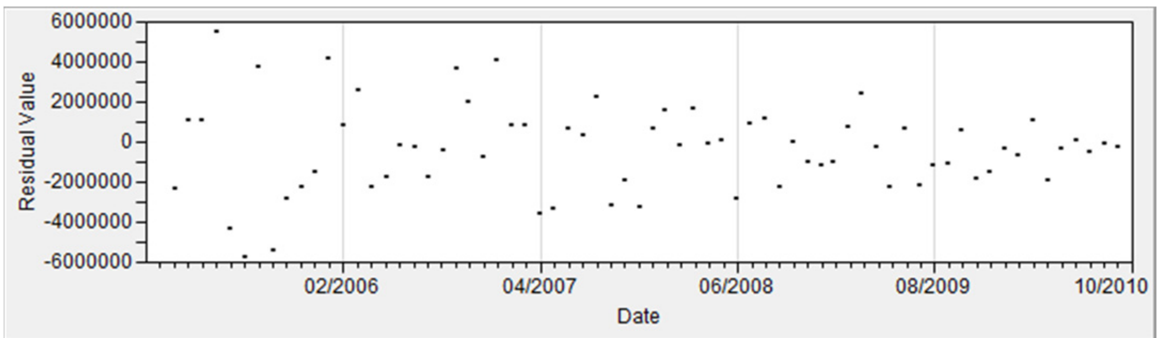

Figure 19: Residuals Plot for SES, Monthly Iraq Demand

The residuals chart for the simple exponential smoothing model is shown in Figure 19. There is a narrowing tendency in the residuals plot attributable to the stabilizing of demand as the Iraq conflict nears its conclusion, but overall the graph shows a normally distributed plot of residuals centered on zero. The resulting fit and forecast graph is given in Figure 20.

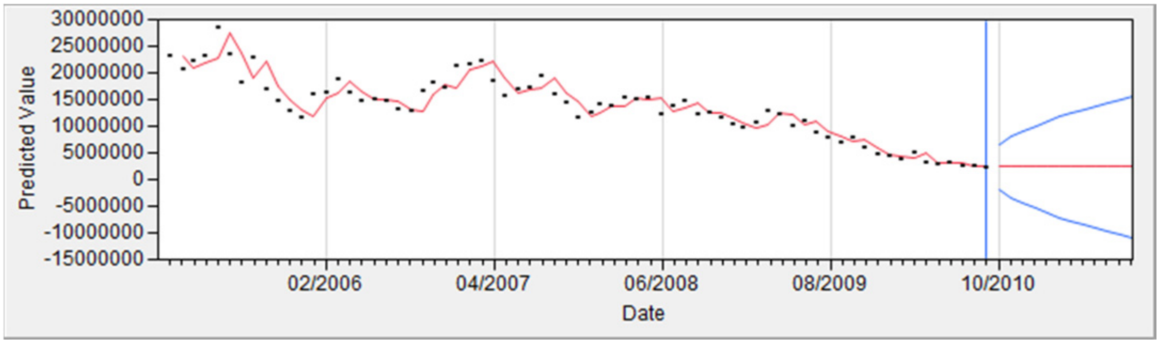

Figure 20: SES Forecast, Monthly Iraq Demand

Model Validation

Table 8: Validation Statistics for Simple Exponential Smoothing Model

RMSEv=	339005.377	RE=	0.999140262
RMSEc=	2203046.962	R2=	0.859
% Difference	84.61	% Difference	16.31

Table 8 shows the validation statistics for the chosen simple exponential smoothing model. Again the RMSE values are not roughly equivalent, supporting previous observations about the Iraq data. However, the validation period shows an explainable *increased* accuracy, not a failure in the forecast model. The RE and R2 values are closer in value, with RE almost equal to one. Therefore, the model is considered validated.

Model Equation

The model equation is given by Equation 3. The smoothing constant as computed by JMP is: α = 0.8591, with a t-ratio of 6.53.

Afghanistan+ Model

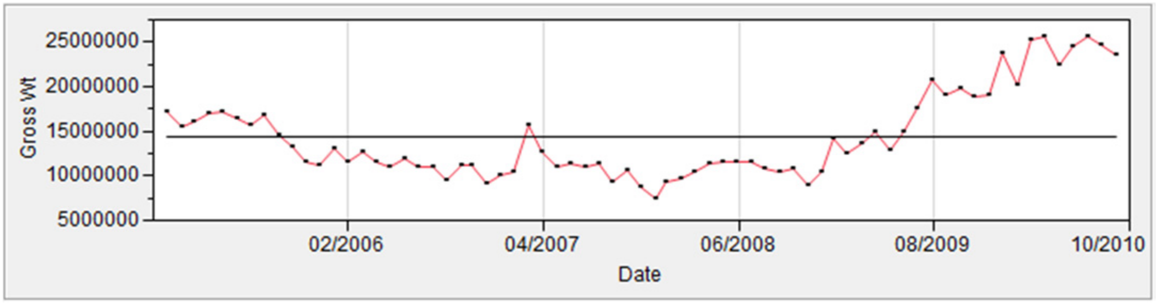

Figure 21: Time Series of Monthly Afghanistan+ Demand, Jan 2005 – Sep 2010

Monthly Afghanistan+ demand from January 2005 through September 2010 is shown graphically in Figure 21. The demand curve saw a reduction from the beginning of 2005 through 2009, likely due to sharing resources with Iraq operations. Post-Iraq demand is definitely higher, corresponding with the increase in boots-on-ground in Afghanistan from Figure 1. A slight seasonal pattern with a 12-month cycle also exists.

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Simple Exponential Smoothing	67	3.072e+12	2149.3057	2151.5252	0.868	2147.3057	6	6	9.563809	1327679.6
Double (Brown) Exponential Smoothing	66	3.105e+12	2121.3150	2123.5196	0.864	2119.315	4	4	10.140919	1400539.8
Linear (Holt) Exponential Smoothing	65	3.093e+12	2122.1188	2126.5282	0.866	2118.1188	5	5	9.811185	1363944.3
Seasonal Exponential Smoothing	54	2.384e+12	1767.8745	1771.9252	0.898	1763.8745	3	2	9.492459	1290944.6
Winters Method (Additive)	53	2.306e+12	1767.0136	1773.0897	0.903	1761.0136	2	3	9.216352	1234474.6
IMA(1, 1)	66	3.091e+12	2150.7088	2155.1478	0.869	2146.7088	7	7	9.604402	1324792.2
Seasonal ARIMA(0, 1, 1)(0, 1, 1)12	53	1.819e+12	1763.6322	1769.7082	0.909	1757.6322	1	1	9.089125	1211634.6

Figure 22: Model Comparison, Monthly Afghanistan+ Demand

The JMP model comparison output for the exponential smoothing and ARIMA techniques is presented in Figure 22. The AIC and SBC rankings agree on seasonal ARIMA as the best model. The RSquare and MAPE values for this method are also the best of the group. Thus the goodness-of-fit measures all point to Seasonal ARIMA as the best model.

The forecast graphs in Appendix H again show that the non-seasonal methods are not likely candidates for Afghanistan+ fields, which exhibit some seasonality. The seasonal ARIMA model, whose fit characteristics were clearly superior, has a forecast that increases more than might be expected for sustainment operations which have apparently reached a new steady state.

Model Selection

Model validation is required for the seasonal ARIMA model to see if it falls victim to a leveling off of demand at the end of the data window.

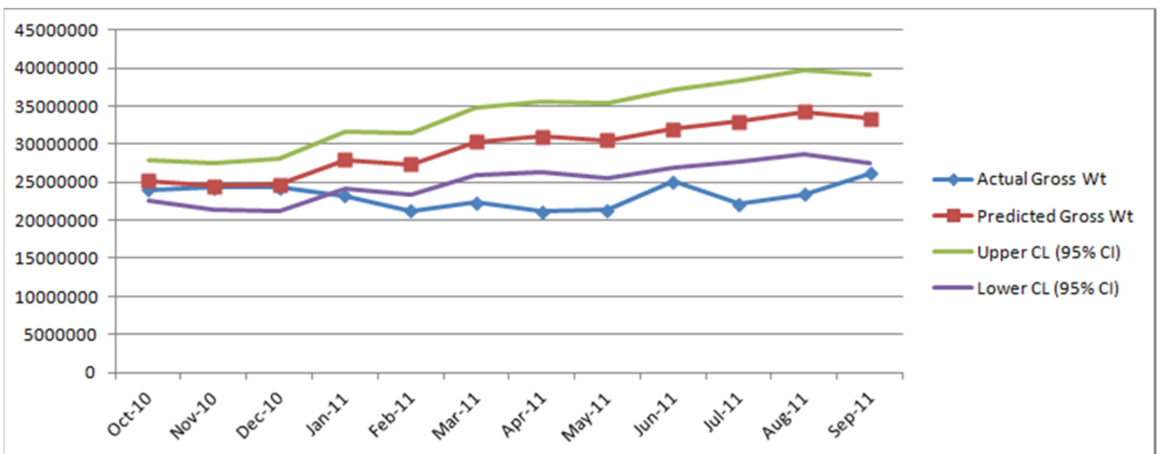

Figure 23: ARIMA(0,1,1)(0,1,1) Forecast, Monthly Aghanistan+ Demand

Looking at Figure 23 the SARIMA model continues to predict increased demand, but the actual demand has leveled off, causing it to quickly depart from the lower control limit of the forecast. The seasonal exponential smoothing model, also with competitive goodness-of-fit criteria in Figure 22, shows a forecast graph in Appendix H that levels off. Figure 24 plots its one-year ahead forecast versus the actual demand. Based on this comparison, the researcher chooses the seasonal exponential smoothing method to model monthly Afghanistan+ demand.

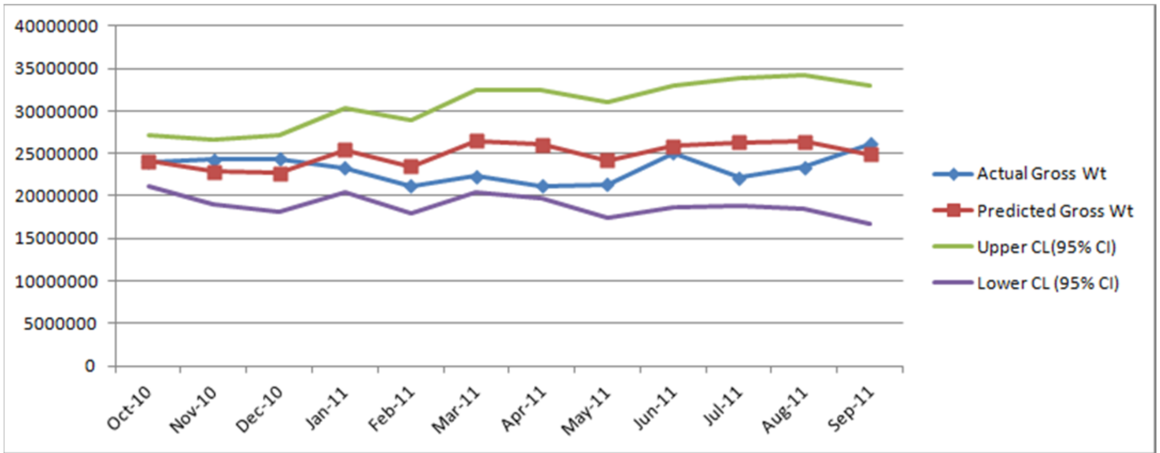

Figure 24: Seasonal Exponential Smoothing, Monthly Afghanistan+ Demand

The residual plot, shown in Figure 25, appears normally distributed with a mean variance close to zero. The model fit and forecast graph is given in Figure 26.

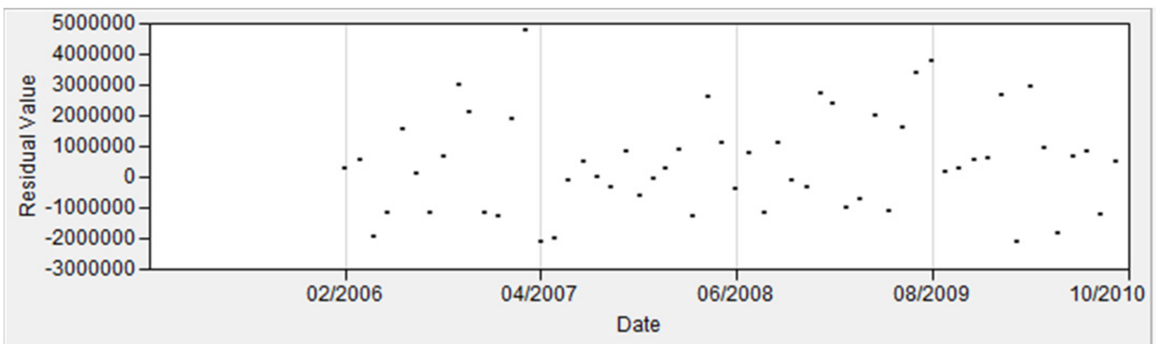

Figure 25: Residuals Plot for Seasonal Exponential Smoothing, Monthly Afghanistan+ Demand

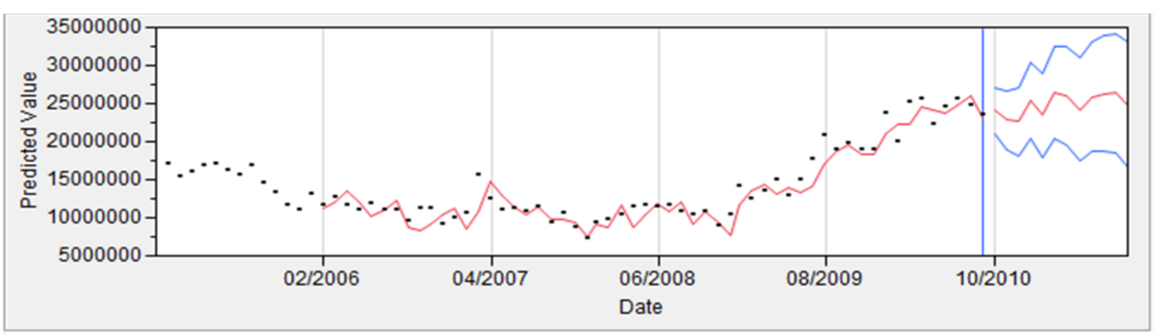

Figure 26: Seasonal Exponential Smoothing, Monthly Afghanistan+ Demand

Model Validation

Table 9: Validation Statistics for Seasonal Exponential Smoothing Model

RMSEv=	2757801.588	RE=	0.918803025
RMSEc=	1651758.082	R2=	0.8
% Difference	66.96	% Difference	14.85

Table 9 shows the validation statistics for the chosen seasonal exponential smoothing model. Again, it is arguable whether the RMSE values are roughly equal, but as with the weekly Afghanistan+ model there is a clearer picture in the RE and R2 values. The RE value is close to one and very close to R2, and therefore the model is considered validated.

Model Equation

The model equation is given by Equations 6.1 – 6.3. The smoothing constants as computed by JMP are:

α = 0.7577, with a t-ratio of 5.59.
δ = 1.000, with a t-ratio of 1.10.

The low t-ratio for the seasonality constant (less than two) corresponds to a 27.5% chance that this smoothing constant is not representative, that is, the seasonality is weak enough that the confidence in the calculated value for the seasonality constant is lower than desired. Nonetheless, there is some degree of seasonality in the model.

Summary

The weekly and monthly demand data for Iraq and Afghanistan+ airfields were analyzed for this research. Thus four distinct models were created and validated in this section. The researcher applied seven time series techniques to each situation, selecting the most promising model from each for validation. Model selection was based on a blended analysis of goodness-of-fit measures and graphical representation of the forecast model. Each selected model was then validated to assess its performance on withheld data.

Both Afghanistan+ groupings were best fit by a seasonal exponential smoothing model, while both Iraq groupings were best fit by a simple exponential smoothing model. In the case of the Iraq models, it is important to note that the end of operations in Iraq has left very little sustainment demand in that AOR. Therefore, the model that was selected as the best fit moving forward would likely not have been the best fit during the campaign.

Finally, there were mixed results as to which grouping size performed better. The weekly groupings in Iraq had higher RSquare scores for the model fit to historical data. However, in the validation period the monthly groupings generally proved better than the weekly groupings, as both the Afghanistan+ and Iraq monthly models scored higher RE than the weekly models. However, in validation the Afghanistan+ monthly model had a higher deviation in RMSE between validation and historical data periods than did the Afghanistan+ weekly model.

V. Conclusions and Recommendations

Chapter Overview

This research addresses a void that exists in the creation of a useful and adaptable forecasting tool for demand of airlifted sustainment cargo. Having described the problem before the CRAF and military planners, the researcher proposes time series forecasting as a solution to this problem. This section provides the "so-what" of this research effort, offering conclusions on the created models. It suggests actions to take in the employment of these models as well as future research that might further improve the models themselves. The paper then concludes with a summary.

Conclusions of Research

The primary objective of this research is to forecast airlift sustainment cargo demand up to one year out in the USCENTCOM AORs. By disaggregating the theater sustainment cargo into AOR-specific demand and grouping the data into weekly and monthly totals, the researcher finds that time series forecast techniques can create suitable models with high RSquare values. The models were validated using reserved data, indicating the models were not merely fit to historical data but also show the ability to forecast future data. In fact, when compared to the five-year moving average technique recently proposed to increase CRAF fixed buys, the researcher's model achieves lower sum of squared error values.

The researcher modeled both Iraq and Afghanistan+ airfields. In the case of the Iraq models, the intent was less to create a predictive model (as Iraq operations have come to an end) as it was to show that historical demand could be fit well using time series techniques. Modeling Iraq also revealed that both AORs, throughout the course of active operations, are best modeled by seasonal exponential smoothing (and possibly seasonal ARIMA, timeframe dependent). This suggests that future conflicts in other AORs might also be fit well using these time series forecasting techniques.

Significance of Research

Time series forecasting has been studied and applied to myriad prediction problems over a spectrum of scenarios, including demand for aviation passenger and cargo transportation. It has on at least one occasion been applied to the specific problem of concern in this paper, namely the airlift of sustainment cargo to overseas military destinations. However, the researcher was not able to find a recent application of modern time series forecasting techniques to the creation of an adaptable forecast model for this problem. This research fills that gap. The models created in this research are immediately useful to the forecasting of both longer term (one year) and shorter term demand. Given the emphasis the CRAF participants put on accurate demand forecasting, which helps their business planning in a competitive environment, any increase in either forecast range or forecast accuracy is welcome. The comparison in Appendix B of the recently proposed moving average method to the Afghanistan+ model from this research indicates the latter gives a much improved forecast judging by demand error over the 2010 and 2011 test years. Thus the DoD could have more confidence in its forecast technique, and therefore could offer a higher fixed buy than with competing methods. CRAF partners and military schedulers would benefit by the increased accuracy.

Recommendations for Action

It is recommended that USTRANSCOM planners take the Afghanistan+ models and use them in parallel with current methods for a number of months. This test period would afford the opportunity to learn how to reduce new data and add it to the model databases, to navigate a data analysis program such as JMP, and to further validate the models by testing them against future months. Assuming the models continue to perform well, the models should be employed as appropriate in the forecasting of short and long-term demand. The strength of an adaptive model is its ability to be updated with new data, yet retain its basic structure. As the forecast graphs in this paper illustrate the forecast confidence intervals can widen fairly quickly. Therefore updates, perhaps monthly, would be crucial to maintaining prediction accuracy.

Recommendations for Future Research

If the GATES database had included unit type data, the researcher would have tried a regression analysis using location, date, and unit type as predictors of demand. Unfortunately unit type data was not immediately available, although the "ultimate consignee" column does list the Department of Defense Activity Address Code (DoDAAC) for the requesting unit. The DoDAAC is an address unique to each unit (but not unit type), so in theory an extensive mapping effort could be accomplished to decode the DoDAACs and uncover the underlying unit type requesting the cargo. Future research to do just that would allow for regression analysis on the same data set, possibly increasing model accuracy.

As alluded to in the literature review, some researchers have attempted to find hybrid approaches to forecasting in order to improve on the accuracy possible to a single method. In that vein, research that fuses together time series forecasting with causal factors, such as boots-on-ground, could conceivably yield better accuracy. An example might include merging a regression analysis, perhaps using boots-on-ground and the variables mentioned above, with a time-series forecast, assigning each method a relative weight. The number of possible hybrid approaches is not limited; there is much potential for model improvement as a result.

Summary

The researcher, using time series forecasting techniques developed primarily in the 1960s and1970s, formed models to forecast sustainment demand in the Iraq and Afghanistan AORs. These models are adaptive, meaning they can be updated with relative ease using newly acquired data. If further validation of the model proves its accuracy, the DoD could with improved confidence increase the CRAF fixed buy. CRAF participants stand to benefit from the larger guaranteed level of business, while DoD schedulers could base their short term sustainment scheduling on the forecast numbers.

Appendix A

Table 10: Aerial Port Codes Represented in the Model

APC	Airfield Name	APC	Airfield Name
3OR	Al Asad, Iraq	KBL	Kabul Int'l, Afghanistan
6OR	Tikrit East, Iraq	KDH	Kandahar, Afghanistan
ADA	Incirlik, Turkey	KIK	Kirkuk, Iraq
ADJ	Marka Int'l, Jordan	KWI	Kuwait Int'l, Kuwait
BAH	Bahrain Int'l, Bahrain	O2R	Al Sahra, Iraq
BSR	Basrah, Iraq	O6R	Qayyarah West, Iraq
DHF	Al Dhafra, UAE	O8R	Tallafar, Iraq
ESB	Esenboga, Turkey	OA1	Bagram, Afghanistan
FAH	Farah, Afghanistan	OA4	Salerno, Afghanistan
FJR	Fujairah, UAE	OR5	Al Taqaddum, Iraq
FRU	Manas, Kyrgyzstan	OR7	Ubayduh Bin Al Jarrah, Iraq
IGL	Cigli, Turkey	OR9	Balad, Iraq
ISB	Islamabad, Pakistan	OSM	Mosul, Iraq
IUD	Al Udeid, Qatar	SDA	Baghdad Int'l, Iraq
JAA	Jalalabad, Afghanistan	TA8	Ali Base, Iraq
JIB	Anbouli, Djibouti	TTH	Thumrait, Oman

Appendix B

The following is a comparison of the seasonal exponential smoothing model and a proposed CRAF incentive model (5-yr moving window), applied to the weekly grouping for Afghanistan+ demand. The moving average was computed using a uniformly weighted moving average (UWMA) with KSigma = 3, and an exponentially weighted moving average (EWMA) with KSigma = 3 and weighting = 0.5.

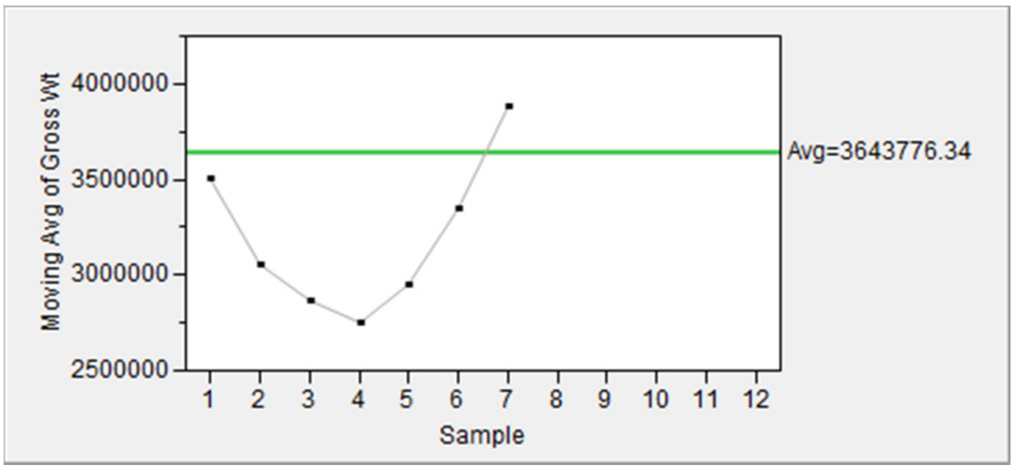

Figure 27: Afghanistan+ Weekly Grouping, 5-Yr UWMA

Table 11: 5-yr UWMA Output

3/0 Cols 7/0 Rows	Sample	n	Moving Avg of Gross Wt
1	1	52	3508865.69
2	2	52	3051013.5
3	3	52	2865471.69
4	4	52	2754910.28
5	5	52	2949325.02
6	6	52	3348542
7	7	52	3880881.48

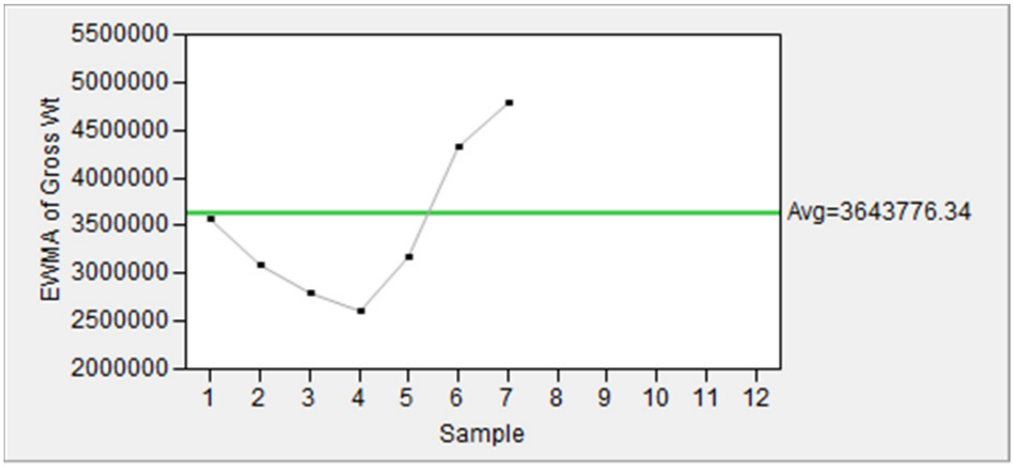

Figure 28: Afghanistan+ Weekly Grouping, 5-Yr EWMA:

Table 12: 5-yr EWMA Output

3/0 Cols / 7/0 Rows	Sample	n	EWMA of Gross Wt
1	1	52	3576321.02
2	2	52	3084741.16
3	3	52	2789564.62
4	4	52	2606395.34
5	5	52	3166689.65
6	6	52	4335820.12
7	7	52	4795339.42

The purpose of this comparison is to determine whether the researcher's model is better for estimating the yearly fixed buy than the approach suggested in the DoD's FY08 budget request—a 5-year moving average of demand for CRAF business (Arthur: 1). The seasonal exponential smoothing model was formed on CY2005-2009 data, and the residuals were captured for the CY2010 forecast. The model was reformed on CY2005-2010 data and again the residuals for the next yearly forecast were captured. The sum-of-squared errors (SSE) for the two periods were computed separately. For the moving average models, the SSE were computed separately for the two forecast periods, CY2010 and CY2011.

The chart below summarizes the results, confirming that the seasonal exponential smoothing model is indeed better than the moving average models in minimizing the SSE. The DoD's proposal was not implemented, although future efforts to increase forecast confidence in order to raise the fixed buy total would be better served by using this research model than the moving average approach.

Table 13: Comparison of Moving Average and Seasonal Exponential Smoothing Methods

	5-Yr UWMA	5-Yr EWMA	Seasonal Exponential Smoothing
Calendar Year	Sum of Squared Errors		
2010	36.31 E+13	30.78 E+13	5.296 E+13
2011	20.94 E+13	64.18 E+13	3.977 E+13

Appendix C

Table 14: Example GATES Database

1	ATMS_ID	PAL_ID	PAL_DT	PLT_GROSS_WT	APC	APOE_MSN_ID	APOE_LEG_ID	MNFST_APOD	APOD_MSN_ID	APOD_LEG	DEP_DT_TM	MDS	TAIL_NUM	MSN_PRIOR	ARR_DT_TM
2	TM6110901815	DHFBMJ	01-Oct-11	986	DHF	6JW52G5XG274	100	KWI	6BW52G6XG274	200	02-Oct-11	KC010A	20192	1B1	02-Oct-11
3	AU1110900025	101TVB	01-Oct-11	4480	DOV	BBW6DP30K27 7	100	3OR	BVW6DP30K27 8	300	04-Oct-11	B74740	N742CK	1B1	04-Oct-11
4	AU1110700501	101TVC	01-Oct-11	1660	WRI	BBW1DY5BL277	100	IUD	BBW1DY60C277	300	04-Oct-11	B74740	N459MC	1B1	05-Oct-11
5	TM5110951143	IUDGPG	01-Oct-11	3680	IUD	AJYF150YT275	100	DHF	AMYF150YT275	200	02-Oct-11	C017A	23293	1B1	02-Oct-11
6	TB4111088580	OA1DKX	01-Oct-11	5800	IUD	AJRA7813A281	200	TTH	AJRA7813A281	300	08-Oct-11	C017A	33124	1B1	08-Oct-11
7	AU1110900108	101TVF	01-Oct-11	5080	WRI	BBWKDG50D27 8	100	KWI	BBWKDG60D27 9	300	05-Oct-11	B74720	N485EV	1B1	05-Oct-11
8	TB4111088580	OA1DKY	01-Oct-11	5800	IUD	AJRA7813A281	200	TTH	AJRA7813A281	300	08-Oct-11	C017A	33124	1B1	08-Oct-11

1	ARR_DT_TM	APOE_APC	APOD_APC	APOE_ICAO	APOD_ICAO	AIR_DIM_CD	PAL_APOE	PAL_APOD	PLT_VOL	PLT_HT	PLT_NET_WT	PLT_PCS_QY	PLT_ULTMT_CNSGNE	PLT_TY_CD	PALLET_TYPE	EVQ_PAL_PS
2	02-Oct-11	DHF	KWI	OMAM	OKBK	D	DHF	KWI	83	40	636	2	FB5819	B	PC	10
3	04-Oct-11	DOV	3OR	KDOV	ORAA	D	DOV	OT4	430	86	4165	1	W91ZKW	L	PC	10
4	05-Oct-11	WRI	IUD	KWRI	OTBH	D	WRI	DHF	486	96	1305	1	W90V9R	L	PC	10
5	02-Oct-11	IUD	DHF	OTBH	OMAM	D	IUD	DHF	378	89	3178	13	FB4811	L	PC	10
6	08-Oct-11	IUD	TTH	OTBH	OOTH	D	OA1	TTH	302	55	5500	1	FB4668	B	PC	10
7	05-Oct-11	WRI	KWI	KWRI	OKBK	D	WRI	KWI	435	86	4725	1	W91RH2	L	PC	10
8	08-Oct-11	IUD	TTH	OTBH	OOTH	D	OA1	TTH	302	55	5500	1	FB4668	B	PC	10

Appendix D

Duke University's "Rules" for Identifying ARIMA Models

Rule 1: If the series has positive autocorrelations out to a high number of lags, then it probably needs a higher order of differencing.

Rule 2: If the lag-1 autocorrelation is zero or negative, or the autocorrelations are all small and patternless, then the series does not need a higher order of differencing. If the lag-1 autocorrelation is -0.5 or more negative, the series may be over-differenced.

Rule 3: The optimal order of differencing is often the order of differencing at which the standard deviation is lowest.

Rule 4: A model with <u>no</u> orders of differencing assumes that the original series is stationary (mean-reverting). A model with <u>one</u> order of differencing assumes that the original series has a constant average trend (e.g. a random walk or SES-type model, with or without growth). A model with <u>two</u> orders of total differencing assumes that the original series has a time-varying trend (e.g. a random trend or LES-type model).

Rule 5: A model with <u>no</u> orders of differencing normally includes a constant term (which represents the mean of the series). A model with <u>two</u> orders of total differencing normally does <u>not</u> include a constant term. In a model with <u>one</u> order of total differencing, a constant term should be included if the series has a non-zero average trend.

Rule 6: If the <u>PACF</u> of the differenced series displays a sharp cutoff and/or the lag-1 autocorrelation is <u>positive</u>--i.e., if the series appears slightly "under-differenced"--then consider adding an <u>AR</u> term to the model. The lag at which the PACF cuts off is the indicated number of AR terms.

Rule 7: If the <u>ACF</u> of the differenced series displays a sharp cutoff and/or the lag-1 autocorrelation is <u>negative</u>--i.e., if the series appears slightly "over-differenced"--then consider adding an <u>MA</u> term to the model. The lag at which the ACF cuts off is the indicated number of MA terms.

Rule 8: It is possible for an AR term and an MA term to cancel each other's effects, so if a mixed AR-MA model seems to fit the data, also try a model with one fewer AR term and one fewer MA term--particularly if the parameter estimates in the original model require more than 10 iterations to converge.

Rule 9: If there is a unit root in the AR part of the model--i.e., if the sum of the AR coefficients is almost exactly 1--you should reduce the number of AR terms by one and increase the order of differencing by one.

Rule 10: If there is a unit root in the MA part of the model--i.e., if the sum of the MA coefficients is almost exactly 1--you should reduce the number of MA terms by one and reduce the order of differencing by one.

Rule 11: If the long-term forecasts appear erratic or unstable, there may be a unit root in the AR or MA coefficients.

Rule 12: If the series has a strong and consistent seasonal pattern, then you should use an order of seasonal differencing--but never use more than one order of seasonal differencing or more than 2 orders of total differencing (seasonal+nonseasonal).

Rule 13: If the autocorrelation at the seasonal period is positive, consider adding an SAR term to the model. If the autocorrelation at the seasonal period is negative, consider adding an SMA term to the model. Do not mix SAR and SMA terms in the same model, and avoid using more than one of either kind.

Appendix E

Weekly Iraq Demand Models

The following graphs show both the fit to historical data as well as the predicted demand and 95% confidence intervals for each model.

Figure 29: Forecast Method Graphs, Weekly Iraq Demand

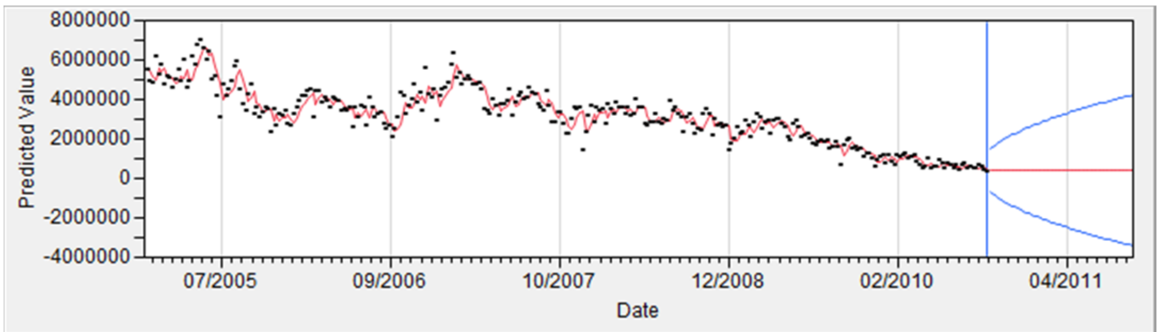

Simple Exponential Smoothing, Weekly Iraq Demand

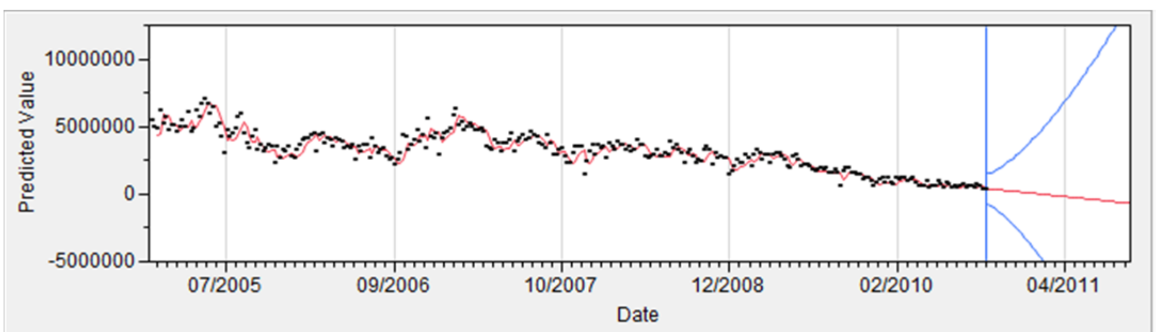

Double (Brown's) Exponential Smoothing, Weekly Iraq Demand

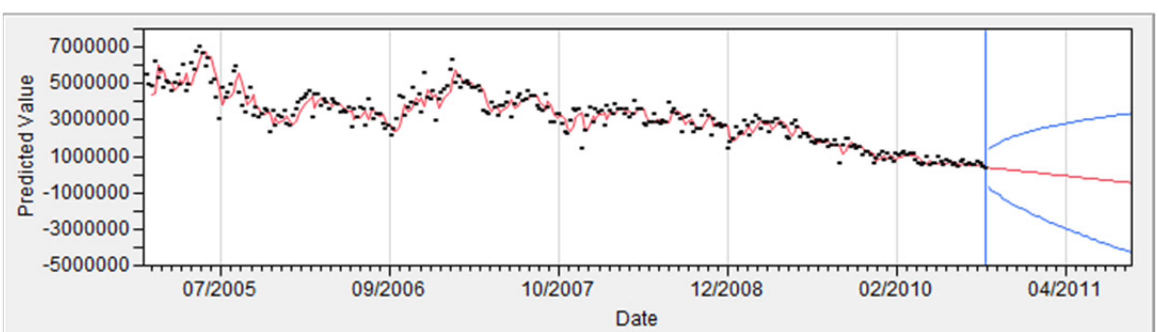

Linear (Holt's) Exponential Smoothing, Weekly Iraq Demand

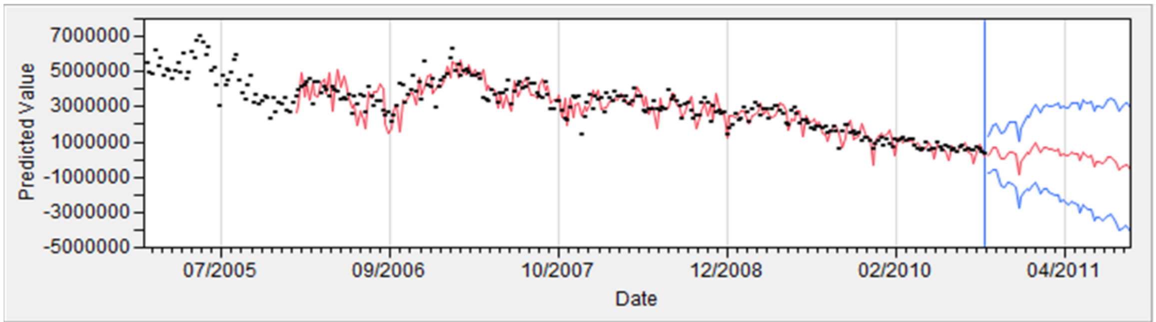

Seasonal Exponential Smoothing, Weekly Iraq Demand

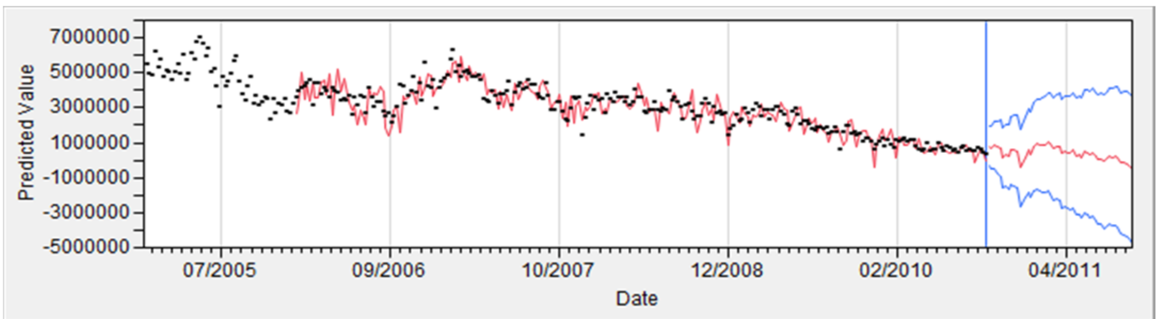

Winters Method, Weekly Iraq Demand

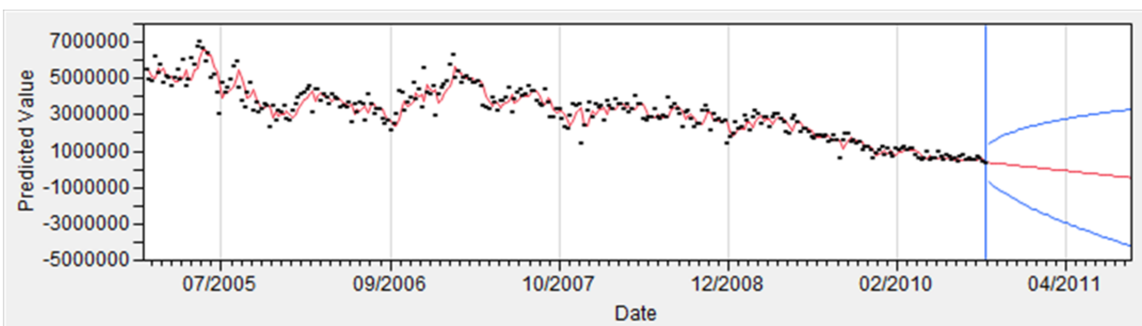

ARIMA(0,1,1) Method, Weekly Iraq Demand

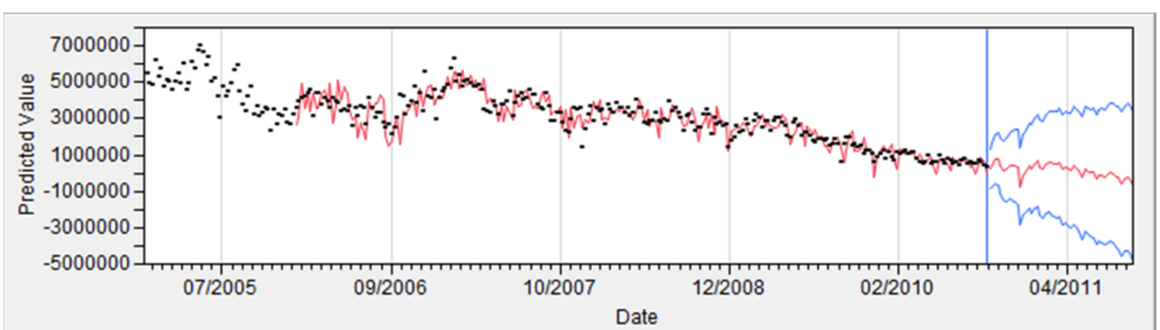

ARIMA(0,1,1)(0,1,1) Method, Weekly Iraq Demand

Appendix F

Weekly Afghanistan+ Demand Models

The following graphs show both the fit to historical data as well as the predicted demand and 95% confidence intervals for each model.

Figure 30: Forecast Method Graphs, Weekly Afghanistan+ Demand

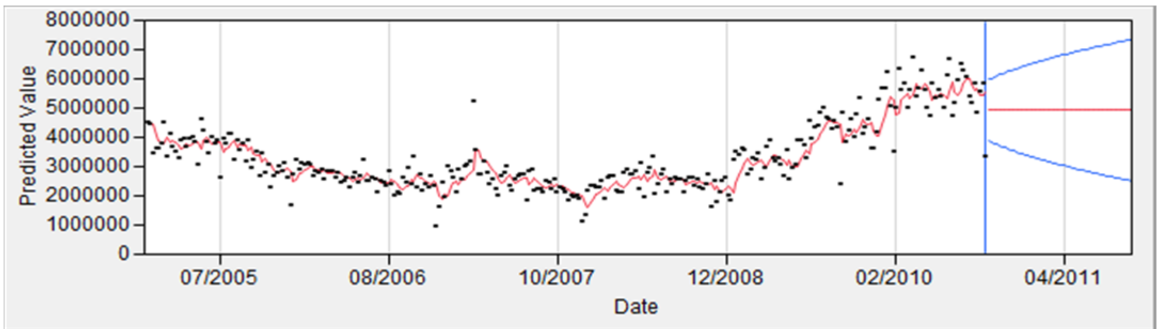

Simple Exponential Smoothing Method, Weekly Afghanistan+ Demand

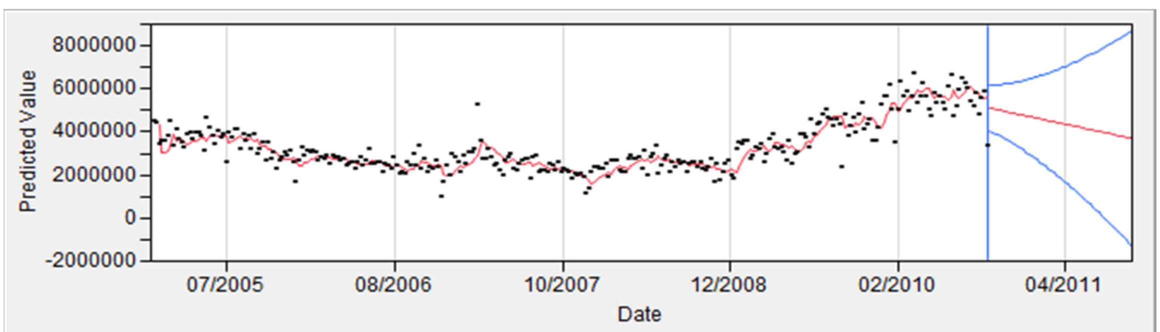

Double (Brown's) Exponential Smoothing Method, Weekly Afghanistan+ Demand

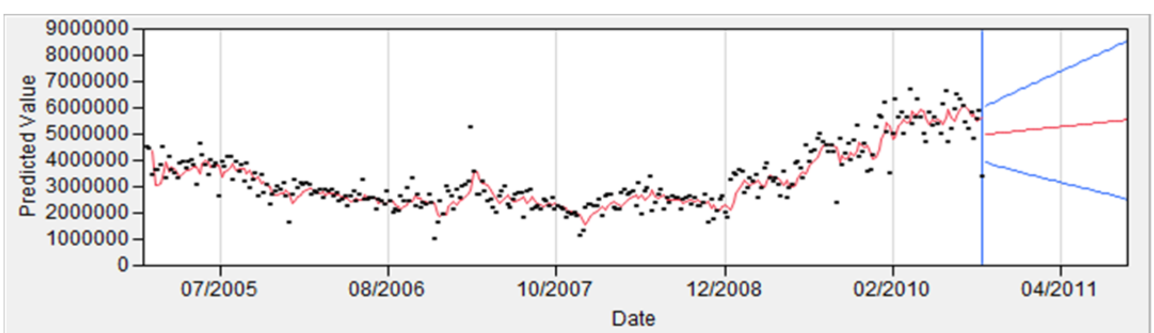

Linear (Holt's) Exponential Smoothing Method, Weekly Afghanistan+ Demand

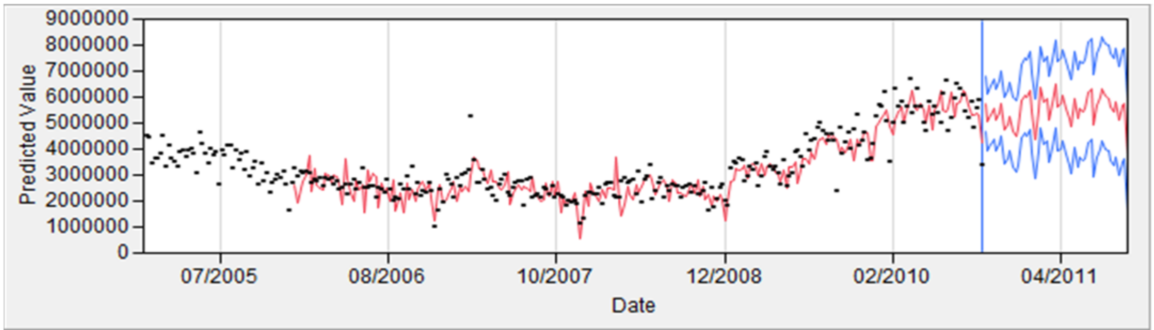

Seasonal Exponential Smoothing Method, Weekly Afghanistan+ Demand

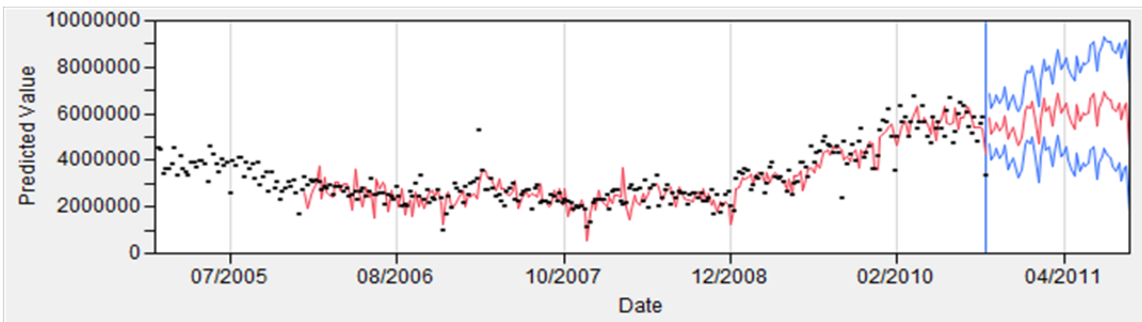

Winters Exponential Smoothing Method, Weekly Afghanistan+ Demand

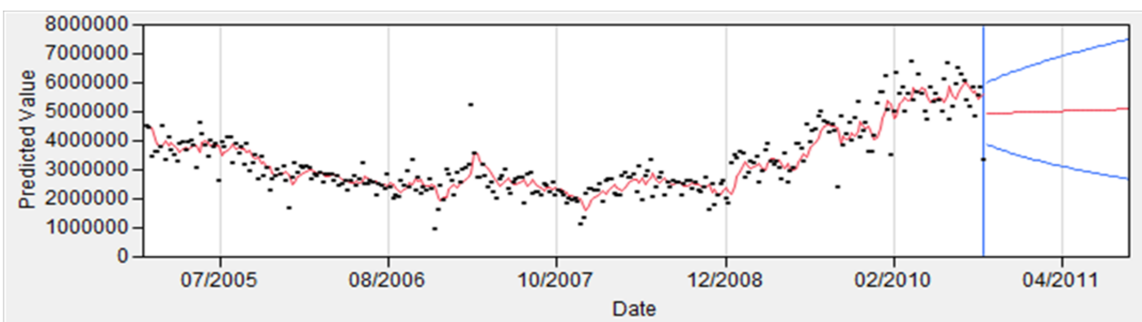

ARIMA(0,1,1) Method, Weekly Afghanistan+ Demand

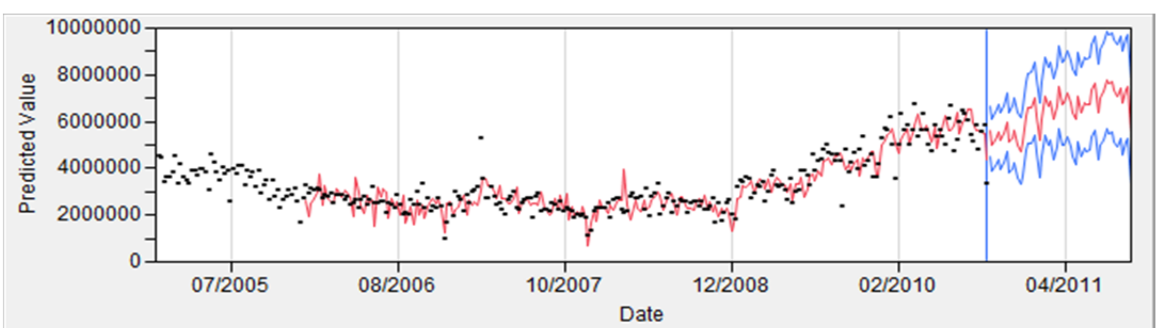

ARIMA(0,1,1)(0,1,1) Method, Weekly Afghanistan+ Demand

Appendix G

Monthly Iraq Demand Models

The following graphs show both the fit to historical data as well as the predicted demand and 95% confidence intervals for each model.

Figure 31: Forecast Method Graphs, Monthly Iraq Demand

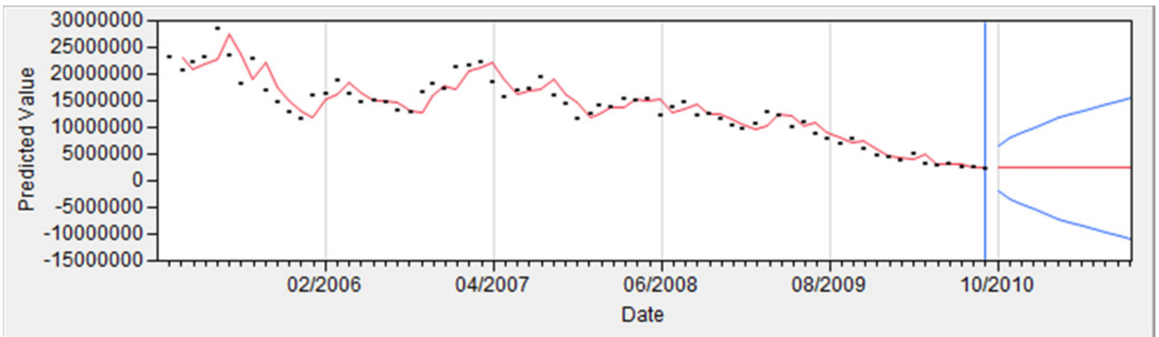

Simple Exponential Smoothing Method, Monthly Iraq Demand

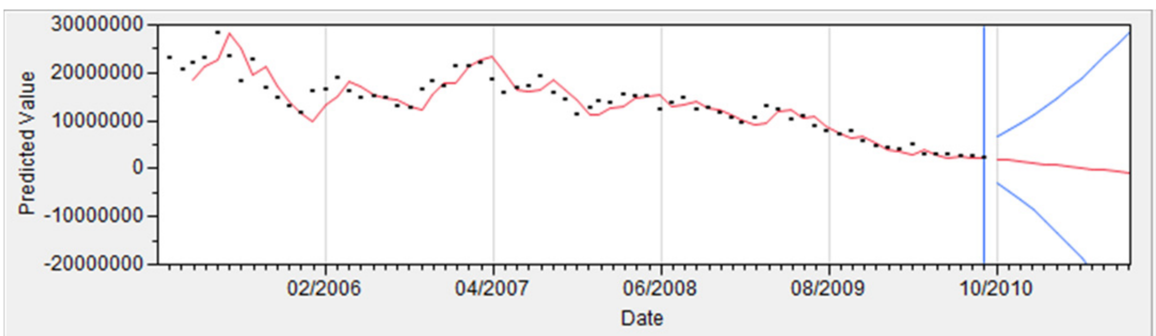

Double (Brown's) Exponential Smoothing, Monthly Iraq Demand

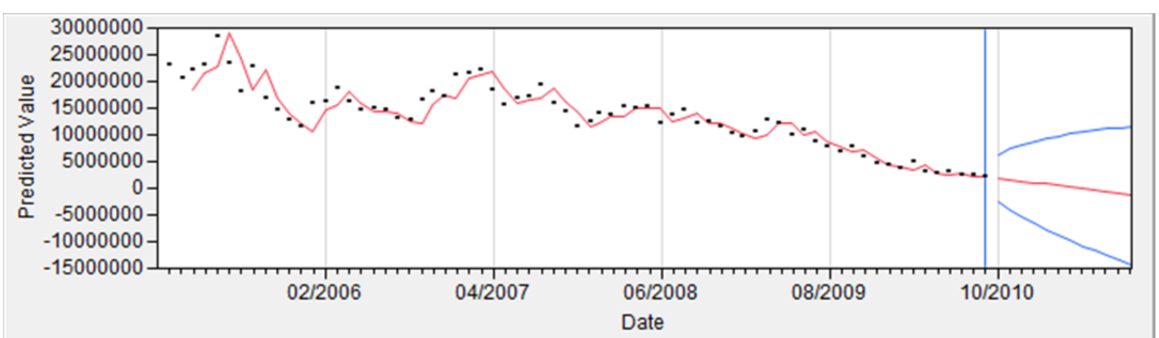

Linear (Holt's) Exponential Smoothing, Monthly Iraq Demand

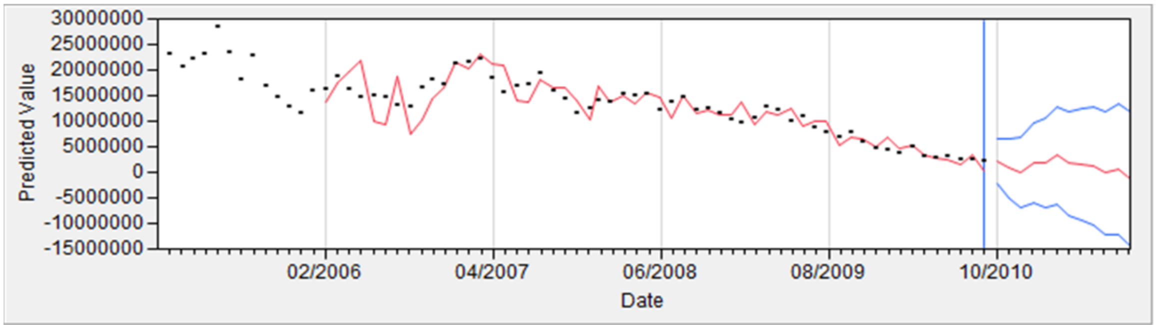

Seasonal Exponential Smoothing, Monthly Iraq Demand

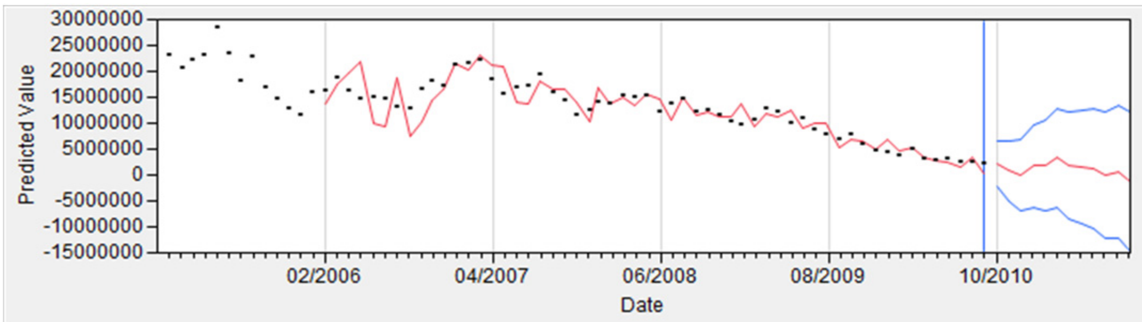

Winters Method, Monthly Iraq Demand

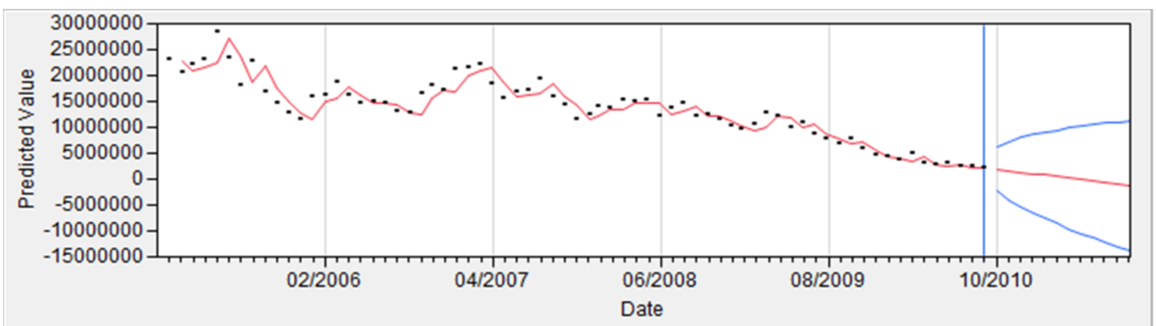

ARIMA(0,1,1), Monthly Iraq Demand

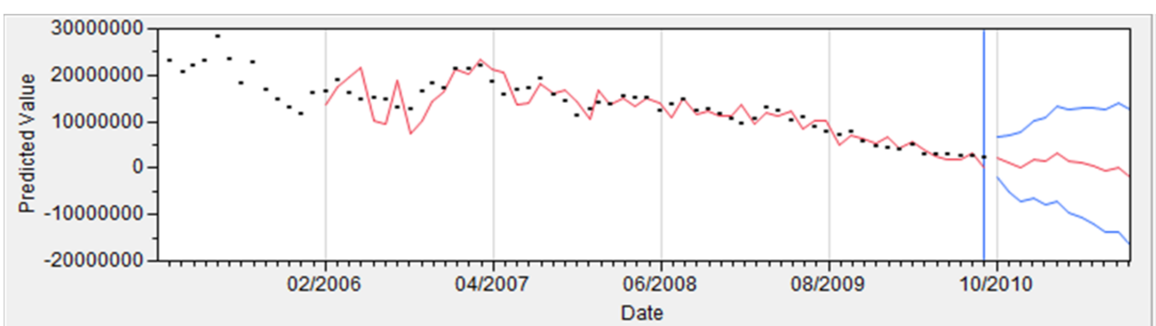

ARIMA(0,1,1)(0,1,1), Monthly Iraq Demand

Appendix H

Monthly Afghanistan+ Demand Models

The following graphs show both the fit to historical data as well as the predicted demand and 95% confidence intervals for each model.

Figure 32: Forecast Method Graphs, Monthly Afghanistan+ Demand

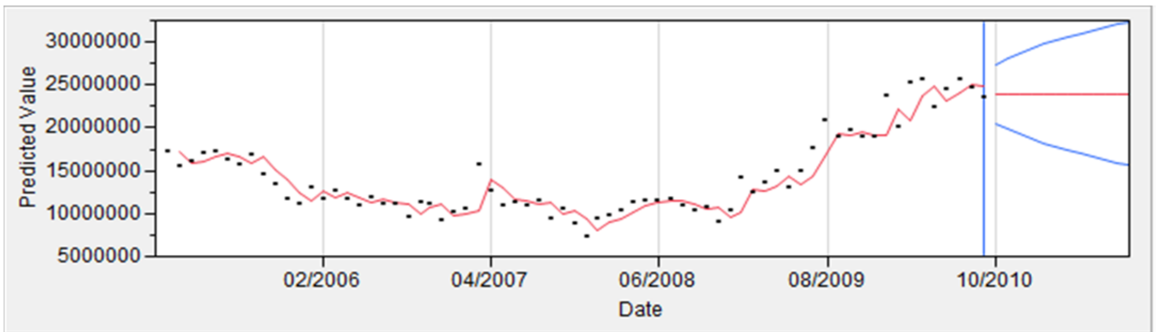

Simple Exponential Smoothing Method, Monthly Afghanistan+ Demand

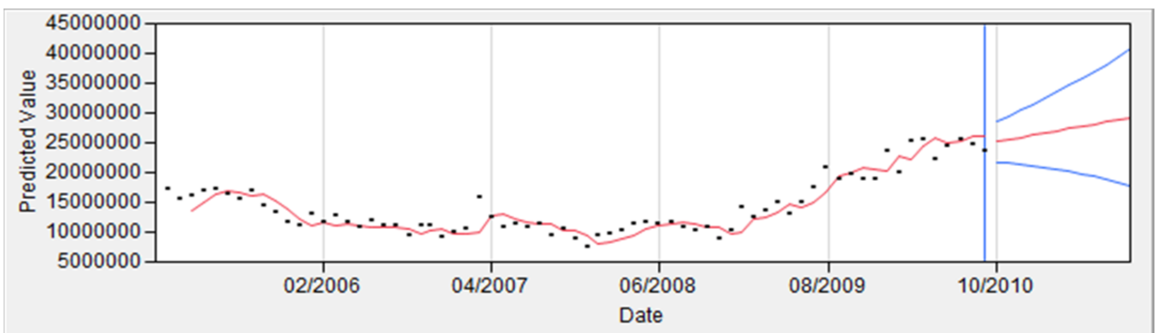

Double (Brown's) Exponential Smoothing Method, Monthly Afghanistan+ Demand

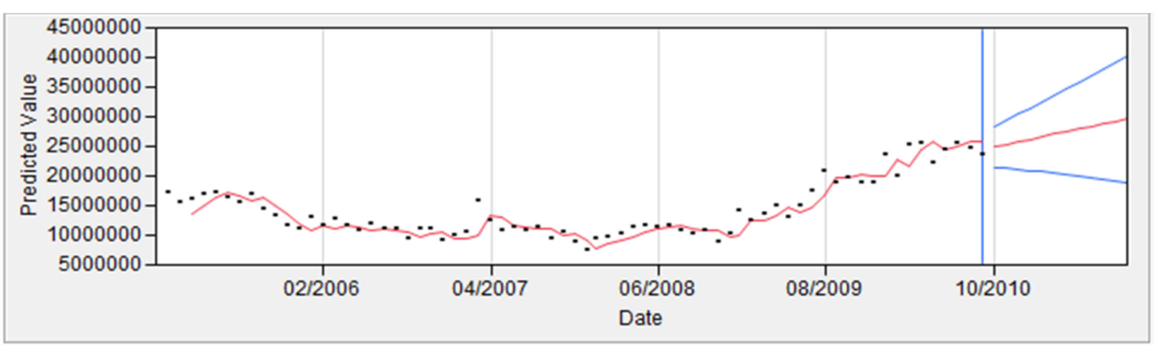

Linear (Holt's) Exponential Smoothing Method, Monthly Afghanistan+ Demand

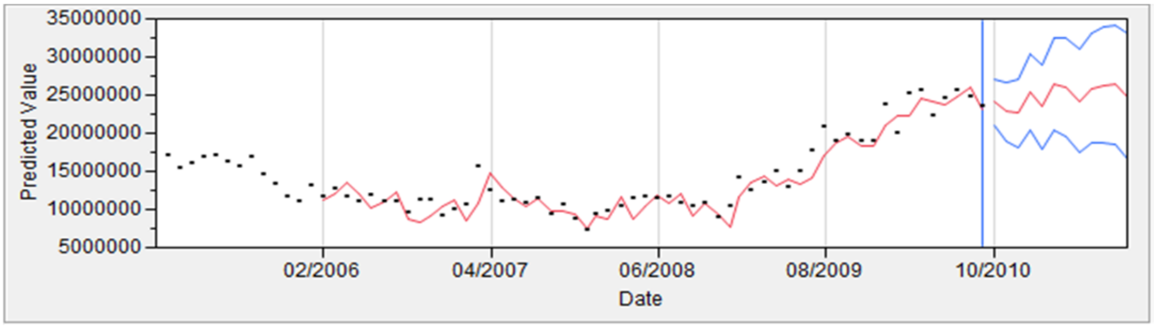

Seasonal Exponential Smoothing Method, Monthly Afghanistan+ Demand

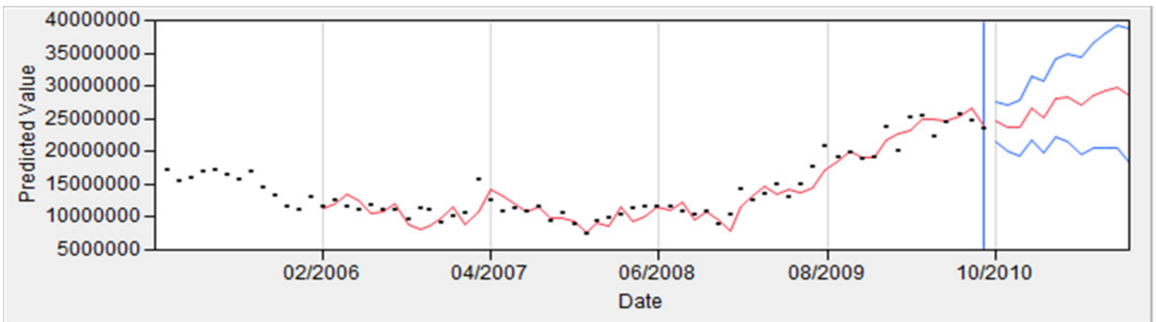

Winters Method, Monthly Afghanistan+ Demand

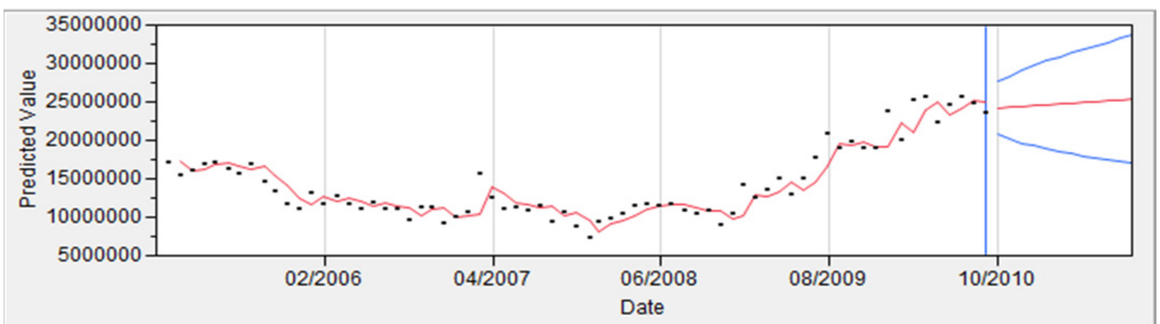

ARIMA(0,1,1), Monthly Afghanistan+ Demand

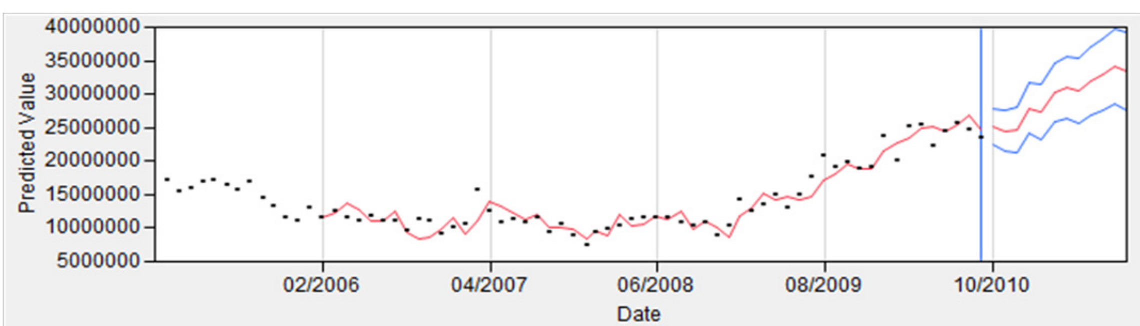

ARIMA(0,1,1)(0,1,1), Monthly Afghanistan+ Demand

Appendix I

This appendix demonstrates the ability of the Afghanistan+ weekly model formed in this paper to be used on an individual airfield, Kandahar, to predict demand over a three month period. This is done in order to address the secondary research objective, namely to create a forecast tool useful to the DoD planners as they schedule traffic flow into and out of individual airfields in the AOR.

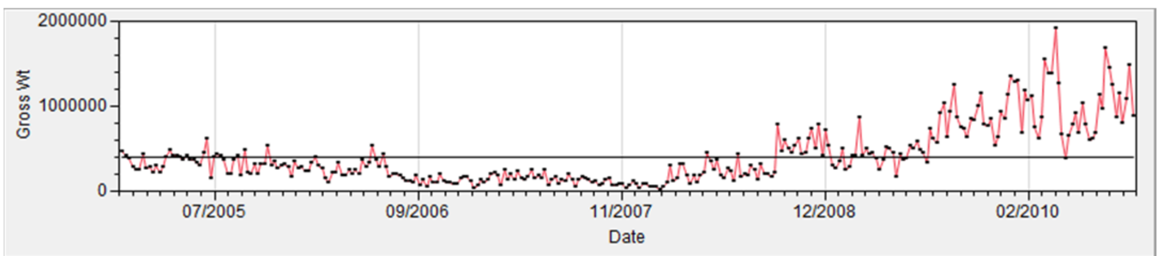

Figure 33: Time Series of Demand at Kandahar Airfield

Weekly Afghanistan+ demand from January 2005 through September 2010 into Kandahar is shown graphically in Figure 33. The demand curve saw an increase from 2008 through 2011. A seasonal pattern is also present, with a 52-week cycle.

The JMP model comparison output for the chosen seasonal exponential smoothing model is presented in Figure 34. The RSquare is again respectable at 0.722:

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Seasonal Exponential Smoothing	245	3.604e+10	6738.0616	6745.0804	0.722	6734.0616	0	0	49.954408	137622.18

Figure 34: Model Comparison Chart for Seasonal Exponential Smoothing, Weekly Kandahar Demand

The final steps in forming the model are a check of the residual plot and a validation of the model on a portion of the data.

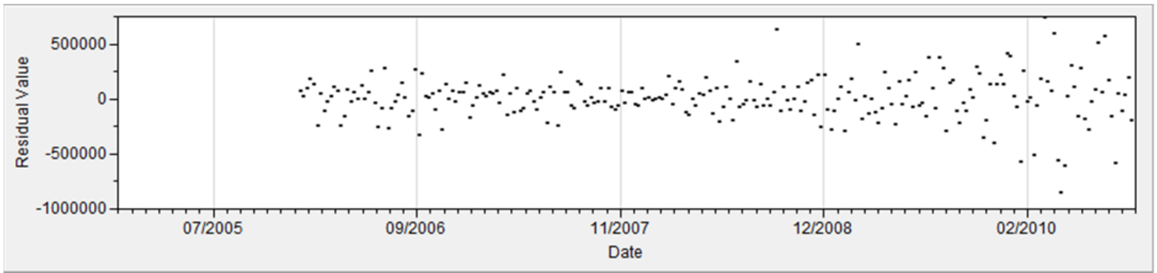

Figure 35: Residuals for Seasonal Exponential Smoothing Model of Kandahar Airfield

The residuals chart for the selected model is shown in Figure 35. The residual plot looks normally distributed, with a mean of zero, although the variance does seem to vary. It is has a period of narrowing at the end of 2007, then disperses some when demand picks up toward the beginning of 2010.

Finally, the model fit and forecast graph is given in Figure 36.

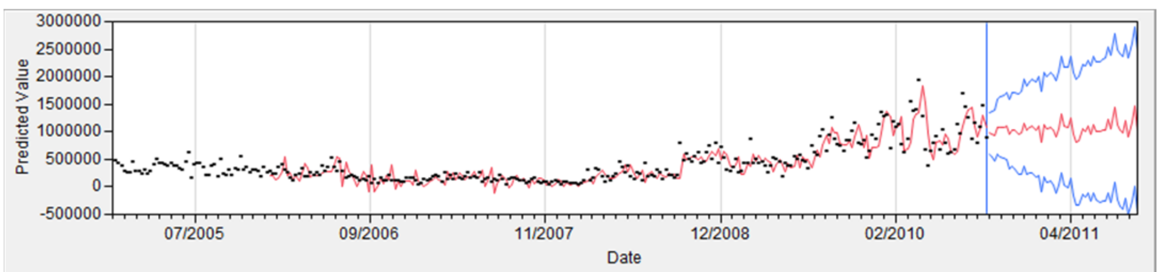

Figure 36: Seasonal Exponential Smoothing Forecast, Kandahar Weekly Demand

Focusing on the forecasted region, the Figure 37 shows the predicted versus actual values for the series over three months. The confidence interval is graphed and the actual data fall within it.

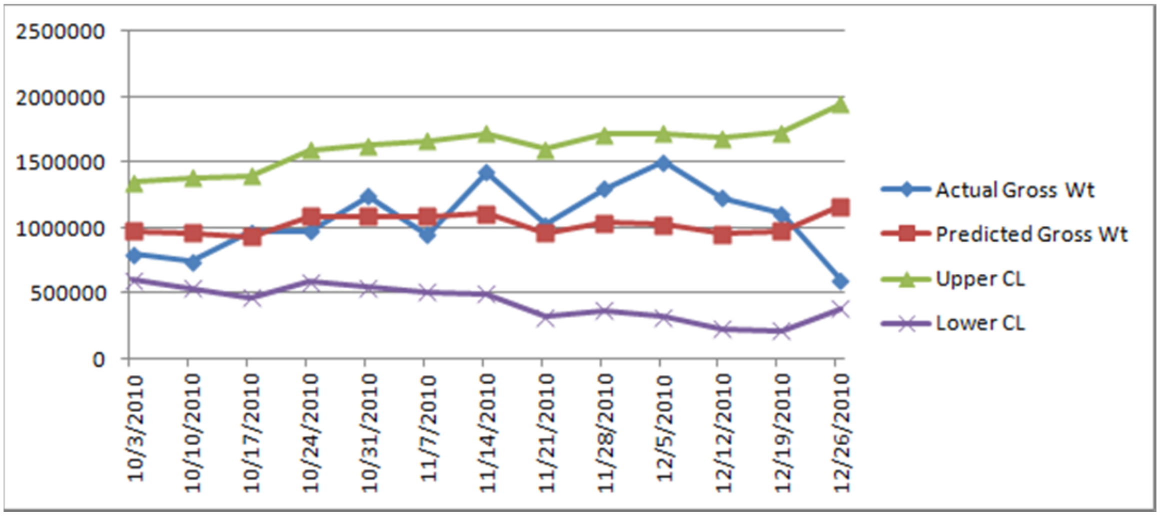

Figure 37: Seasonal Exponential Smoothing 3-Month Forecast, Weekly Kandahar Demand

Model Validation

Table 15: Validation Statistics for Seasonal Exponential Smoothing Model

RMSEv=	269816.2658	RE=	0.822762176
RMSEc=	195878.8829	R2=	0.722
% Difference	37.75	% Difference	-13.96

Table 15 shows the validation statistics for the chosen seasonal exponential smoothing model. The RMSE values are closer than they were for the one-year Afghanistan+ demand, though again it is a subjective matter as to whether they are roughly equal. Yet the RE and R2 values again agree fairly well, and RE is highly positive. Therefore, the model is considered validated.

Model Equation

The model equation is given by Equations 6.1 – 6.3. The smoothing constants as computed by JMP are statistically significant:

α = .5315, with a t-ratio of 7.53.
δ = .8438, with a t-ratio of 4.61.

Bibliography

Air Mobility Command. *MAF Mission ID Encode/Decode Procedures.* Scott AFB, IL: Air Mobility Command, 2009.

Arthur, David. "Issues Regarding the Current and Future Use of the Civil Reserve Air Fleet," Congressional Budget Office, October 2007.

Belasco, Amy. "The Cost of Iraq, Afghanistan, and Other Global War on Terror Operations Since 9/11," Congressional Research Service, 29 March 2011.

Ben-Akiva, M. "Structure of Passenger Travel Demand Models," Massachusetts Institute of Technology, June 1973.

Billion, William E. and Lawrence G. Regan. "A Methodology for Accurately Predicting Demand for Airlift of Military Cargo to Overseas Destinations," presented at the thirteenth National Meeting of the Operations Research Society of America, October 1966.

Box, George and Gwilym Jenkins. *Time series analysis: Forecasting and control,* San Francisco: Holden-Day, 1970.

Chen, Argon, C.-H. Hsu and J. Blue. "Demand planning approaches to aggregating and forecasting interrelated demands for safety stock and backup capacity planning," International Journal of Production Research, May 2007.

"Decision 411 Forecasting," http://www.duke.edu/~rnau/411avg.htm. Accessed 10 April 2012.

"Engineering Statistics Handbook," http://www.itl.nist.gov/div898/handbook/pmc/section4/pmc444.htm. Accessed 12 April 2012.

Garvett, Donald S. and Nawal K. Taneja. "New Directions for Forecasting Air Travel Passenger Demand," Massachusetts Institute of Technology, July 1974.

Graham, David. "Sustaining the Civil Reserve Air Fleet (CRAF) Program," Institute for Defense Analyses, May 2003.

Grismer, Michael W. Jr. "Transforming the Civil Reserve Air Fleet (CRAF) to Enable Combat Power in a Fiscally Constrained Environment," Naval War College Thesis, October 2010.

GEOS 585A Courseware. http://www.ltrr.arizona.edu/~dmeko/notes_12.pdf. Accessed 2 May 2012.

"JMP Support", http://www.jmp.com/support/help/Modeling_Reports.shtml. Accessed 26 April 2012.

JP 3-17, Joint Publication 3-17, Chairman of The Joint Chiefs of Staff (CJCS), *Air Mobility Operations*, 2 October 2009.

JP 4-09, Joint Publication 4-09, Chairman of The Joint Chiefs of Staff (CJCS), *Distribution Operations*, 5 February 2010.

Kress, Moshe. *Operational Logistics: The Art and Science of Sustaining Military Operations,* Boston: Kluwer Academic Publishers, 2002.

Morrison, Steven A. and Clifford Winston. "What's Wrong with the Airline Industry? Diagnosis and Possible Cures," Hearing before the Subcommittee on Aviation Committee on Transportation and Infrastructure United States House of Representatives, 28 September 2005.

Myung, In Jae. "Tutorial on Maximum Likelihood Estimation", Journal of Mathematical Psychology, 16 October 2002.

"STAT 510—Applied Time Series Analysis," https://onlinecourses.science.psu.edu/stat510/?q=node/70. Accessed 11 April 2012.

Stoimenova, E., K. Prodanova and R. Prodanova. "Forecasting Electricity Demand by Time Series Models," American Institute of Physics, 2007.

"The Likelihood Function, Maximum Likelihood Estimator (MLE), Logit & Probit, Count Data Regression Models", http://econweb.rutgers.edu/tsurumi/liklihood.pdf. Accessed 27 April 2012.

Totamane, Raghavendra, Amit Dasgupta, Ravindra N. Mulukutla, and Shrisha Rao. "Air Cargo Demand Prediction," IEEE International Systems Conference, 2009.

TRANSCOM/J3 Operations Update. http://www.ndtascottstlouis.org/Attachments/Luncheons/Presentations/TRANSCOM%20Ops%20Update.pdf/ Accessed 3 March 2012.

U.S. Census Bureau, Statistical Abstract of the United States: 2012: U.S. Scheduled Airline Industry—Summary: 1995 to 2009. http://www.census.gov/compendia/statab/2012/tables/12s1073.pdf Accessed 5 April 2012.

Velonius, Platon M. "Forecasting Tanker Freight Rates," Massachusetts Institute of Technology, 1994.

Wickham, Richard R. "Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation," Massachusetts Institute of Technology, June 1995.

Time Series Forecasting of Airlift Sustainment Cargo Demand

Maj Daniel S. DeYoung
Advisor: Jeff Weir, Ph.D.
Advanced Studies of Air Mobility (ENS)
Air Force Institute of Technology

Introduction

The US has been at war in the SWA AOR since October, 2001. It has had to sustain tens of thousands of troops in both Afghanistan and Iraq operations during that time. The Civil Reserve Air Fleet has been a critical enabler of our strategic airlift of sustainment cargo during these operations. Unfortunately, the airline industry has suffered since the attacks of September 11, 2001, and the difficult economy has put the industry under intense pressure. As a result, it is advantageous for the DoD to strengthen the financial health of its CRAF partners by giving more guaranteed business to them in the annual fixed-buy. This can be made possible through better long-term forecasting of sustainment demand. This research applies time series forecasting techniques to model sustainment cargo demand, using historical FY05-11 data from the GATES database as a source. It models the Iraq and Afghanistan demand signals individually, and competes seven time series techniques to develop the best forecast for each region.

Research Goals

- Understand desire of CRAF partners for better forecasting.
- Quantify historical demand signal from SWA for sustainment cargo
- Model weekly and monthly demand requirement for strategic airlift of sustainment cargo into SWA AOR.

AIR FORCE INSTITUTE of TECHNOLOGY

General Framework

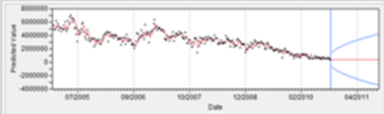

Figure 8: SES Forecast, Weekly Iraq Demand

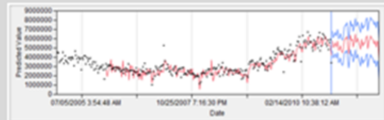

Figure 14: Seasonal Exponential Smoothing Forecast, Weekly Afghanistan+ Demand

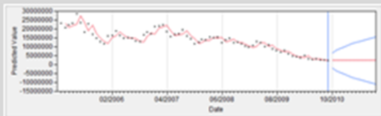

Figure 22: SES Forecast, Monthly Iraq Demand

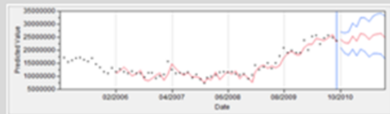

Figure 28: Seasonal Exponential Smoothing, Monthly Afghanistan+ Demand

Model Comparison

Model	DF	Variance	AIC	SBC	RSquare	-2LogLH	AIC Rank	SBC Rank	MAPE	MAE
Simple Exponential Smoothing	67	3.072e+12	2149.3057	2151.5252	0.868	2147.3057	6	6	9.563809	1327679.6
Double (Brown) Exponential Smoothing	66	3.105e+12	2121.3150	2123.5196	0.864	2119.315	4	4	10.140919	1400539.8
Linear (Holt) Exponential Smoothing	65	3.093e+12	2122.1188	2126.5282	0.866	2118.1188	5	5	9.811185	1363944.3
Seasonal Exponential Smoothing	54	2.384e+12	1767.8745	1771.9252	0.898	1763.8745	3	2	9.492459	1290944.6
Winters Method (Additive)	53	2.306e+12	1767.0136	1773.0897	0.903	1761.0136	2	3	9.216352	1234474.6
IMA(1, 1)	66	3.091e+12	2150.7088	2155.1478	0.869	2146.7088	7	7	9.604402	1324792.2
Seasonal ARIMA(0, 1, 1)(0, 1, 1)12	53	1.819e+12	1763.6322	1769.7082	0.909	1757.6322	1	1	9.089125	1211634.6

Figure 24: Model Comparison, Monthly Afghanistan+ Demand

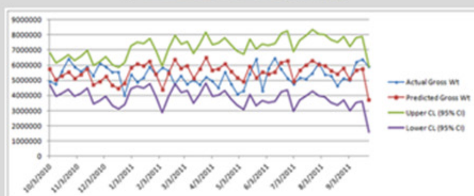

Application

- Seasonal Exponential Smoothing Equation
- Models Afghanistan AOR, Monthly and Weekly Groupings

$$L_t = \alpha(Y_t - S_{t-p}) + (1-\alpha)L_{t-1} \quad (6.1)$$

$$S_t = \delta(Y_t - L_t) + (1-\delta)S_{t-p} \quad (6.2)$$

$$\hat{Y}_t(k) = L_t + S_{t-p+k} \quad (6.3)$$

Where:
L = smoothed value of mean term
S = smoothed value of seasonality term
α = smoothing constant between 0 and 1
δ = seasonality smoothing constant between 0 and 1
$\hat{Y}$ = forecast value

Motivation

- Provide a forecasting tool for use by the DoD to better predict long-term (1-yr) sustainment cargo demand and short-term (3 mos) sustainment cargo demand

Impacts/Contributions

- Better forecasting of long-term demand can enable the DoD to purchase CRAF fixed-buys, helping the civil carriers to manage their business and strengthening the CRAF program.
- Better short-term forecasting can aid DoD planners in scheduling strategic resupply missions and assigning appropriate aircraft types.

Collaboration

USAF AMC/A9

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From – To)*
15 Jun 2012	Master's GRP	23 May 2011 – 15 Jun 2012

4. TITLE AND SUBTITLE
Time Series Forecasting of Airlift Sustainment Cargo Demand

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
DeYoung, Daniel S., Major, USAF

5d. PROJECT NUMBER
N/A

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S)
Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/ENV)
2950 Hobson Way, Building 640
WPAFB OH 45433-8865

8. PERFORMING ORGANIZATION REPORT NUMBER
AFIT/IMO/ENS/12-05

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
USTRANSCOM JDPAC
Mr. Bruce Busler, Director, Joint Distribution Process Analysis Center
1 Soldier Way, HQ SDDC/JDPAC
Scott AFB, IL, 62225
(COMM) (618) 220-5118, (DSN) 770-5118, Bruce.Busler@ustranscom.mil

10. SPONSOR/MONITOR'S ACRONYM(S)
Dir, JDPAC

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A, Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT Forecasting demand for airlift of sustainment cargo is an important function for logistics planners. For the civil reserve air fleet participants (CRAF), having a useful long-range forecast enables them to make business decisions to maximize profit and manage their fleets. Because the DoD relies on CRAF for much of its steady-state and wartime surge requirements, it is important for these civilian enablers to stay financially healthy in what has become a difficult market. In addition to the CRAF airlines, DoD schedulers also benefit from somewhat shorter-term forecasts of demand, as accurate forecasts help them allocate aircraft type and determine route frequency for airlift of sustainment cargo.

Time series forecasting is a method applied in many circumstances, to include forecasting of aviation service demand. It does not require the modeler to attribute causation, but rather uses historical data of a univariate series to predict future values. This paper applies a variety of time-series techniques to historical sustainment demand data from Iraq and Afghanistan AORs, ultimately choosing a single technique to develop into a prediction model for future demand in each AOR. The resulting models show excellent goodness-of-fit values and are successfully validated against a reserved portion of data.

15. SUBJECT TERMS
Forecasting, Sustainment, Airlift, Cargo, Demand, CRAF

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			DeYoung, Daniel S., Maj, USAF
U	U	U	UU or SAR	95	**19b. TELEPHONE NUMBER** *(Include area code)* (609) 754-7740 (daniel.deyoung@us.af.mil)